

Phase Noise

Art, Science and Experimental Methods

Enrico Rubiola

Le macchine del vento, ©Roberto Bergonzo



Besançon, France, May 2023

Download the
Companion paper




Download the
Enrico's Chart



Thanks to FIRST-TF <https://first-tf.com>

Clock signal $v(t) = V_0[1 + \alpha(t)] \cos[2\pi\nu_0 t + \varphi(t)]$



Boldface notation

total = nominal + fluctuation

$\varphi(t) = 2\pi\nu_0 t + \varphi(t)$ phase

$\nu(t) = \nu_0 + (\Delta\nu)(t)$ frequency

$x(t) = t + x(t)$ time

$y(t) = 1 + y(t)$ fractional frequency

Phase noise spectrum

Definition

$S_\varphi(f)$ [rad^2/Hz] is the one-sided PSD ($f > 0$) of $\varphi(t)$

$$S_\varphi(f) = 2\mathcal{F}\{\mathbb{E}\{\varphi(t)\varphi(t+\tau)\}\}, \quad f > 0$$

Evaluation

$$S_\varphi(f) = \frac{2}{T} \langle \Phi_T(f) \Phi_T^*(f) \rangle_m$$

avg on m data, $\Phi_T(f)$ = DFT of $\varphi(t)$ truncated on T

Deprecated: most often, 'phase noise' refers to $\mathcal{L}(f)$

Only $10\log_{10}[\mathcal{L}(f)]$ is used, given in dBc/Hz

Definition: $\mathcal{L}(f) = \frac{1}{2} S_\varphi(f)$ [the unit c/Hz never used]

Literally, the unit 'c' is a squared angle, $\sqrt{c} = \sqrt{2} \text{ rad} \approx 81^\circ$

Two-sample (Allan-like) variances


Definition

$$\sigma_y^2(\tau) = \mathbb{E}\left\{\frac{1}{2}[\bar{y}_2 - \bar{y}_1]^2\right\} \quad y(t) \rightarrow \bar{y} \text{ averaged over } \tau$$



\bar{y}_2 and \bar{y}_1 are contiguous
Bare mean $\bar{y} \rightarrow$ Allan variance AVAR
Weighted averages \rightarrow MVAR, PVAR, etc.

Evaluation

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} [\bar{y}_{k+1} - \bar{y}_k]^2 \quad M \text{ contiguous samples of } \bar{y}$$



Frequency fluctuation PSD \leftrightarrow Allan Variance




DOI: 10.5281/zenodo.4399218 (always latest version)

© Enrico Rubiola
ORCID 0000-0002-5364-1835

Manufacturers of equipment for time & frequency can give the Chart for free with their products. Details on the download URL.

February 10, 2023
Page 1/2




Additive Noise
RF noise added to the carrier
Statistically independent AM & PM


Parametric Noise
Near-dc noise modulates the carrier
AM & PM related and narrowband

Leeson effect


Oscillator's PM-noise equivalent







Time Domain



Frequency Domain



Spectra and Polynomial Law



Phase Noise

Enrico Rubiola

February 7, 2025

The usual page on the back side of the title

Contents

1	Basic Facts	2
1.1	Physical quantities	3
1.2	Spectra	8
1.3	The deprecated quantity $\mathcal{L}(f)$	10
1.4	Two-port components	15
1.5	Scaling the carrier frequency up and down	18
1.6	Phenomenology of the phase-noise processes	23
1.7	Statistics and uncertainty	23
1.8	Amplitude and phase modulation	23
1.9	Suggested readings	25
2	Random Noise	27
2.1	Thermal (Johnson) noise	27
2.2	Shot (Schottky) noise	32
2.3	Quantum noise	34
2.4	Flicker (1/f) noise	37
2.5	Technical representations of noise	42
2.6	Noise in selected devices	48
2.7	The Measurement of noise temperature	48
2.8	Points of interest	50
3	Power Spectra	53
3.1	Introduction	54
3.2	Power spectral density	56
3.3	The cross-spectrum method	58
3.4	Running the experiment	59
3.5	Estimation of S_{xx}	63
3.6	Estimation of S_{yx} and noise rejection	64
3.7	Statistical independence on the frequency axis	73
3.8	Applications and experimental techniques	74
3.9	AppendixMathematical background	79
3.10	Windowing	86
3.11	Overlapping	86
4	Allan variances	87
4.1	Conversions	97
4.2	Confidence Intervals	101
4.3	Software Tools	103

4.4	Stefan Droste	105
5	Two-Port Devices	107
6	Digital Electronics	109
6.1	Introduction	110
6.2	Definitions, and Phase Noise Models	110
6.3	Noise in the Clock Distribution	113
6.4	Selected Noise Measurements	118
6.5	The Volume Law	121
6.6	Input Chatter	123
6.7	Internal PLL	124
6.8	Thermal Effects	126
7	Systems	131
8	Power Splitters	133
8.1	Basic facts	133
8.2	Common power splitters, impedance matched	136
8.3	The weird case of the sub-thermal oscillator	142
8.4	Other effects	143
9	Analog Phase Noise Analyzers	147
10	Digital Phase Noise Analyzers	149
11	Analyzers	151
11.1	Aeroflex	151
11.2	Anapico	151
11.3	Anritsu	151
11.4	BNC	151
11.5	Andrew Holme's hobby project	151
11.6	Holzworth	159
11.7	John Miles and Jackson Labs	159
11.8	Keysight E5052B, optional E5053A	160
11.9	Keysight, the new analyzer	160
11.10	Koheron 250 MSPS acquisition board	160
11.11	Meir Alon's project	161
11.12	Microchip 5120A	161
11.13	Microchip 5125A	161
11.14	John Miles' TimePod	161
11.15	Rohde Schwarz FSWP and FSPN	161
11.16	NoiseXT (Spherea Arcale)	165
11.17	tinyPFA	167
11.18	Vremya VCH-314/315/315M/325	168
11.19	Yanjun Ma BG6KHC's project	168

12 Photonic Systems	171
12.1 Basic concepts	172
12.2 Optical-fiber microwave delay	176
12.3 Dual-delay phase noise measurement	180
12.4 Experiments	183
12.5 Further developments	188
12.6 Photodetector	189
12.7 Common mistakes	191
12.8 High-speed photodetectors	192
12.9 Method and choice of the measurement scheme	192
12.10Background noise	192
12.11Experimental problems	192
12.12Results	192
12.13Photodetector	192
12.14Introduction	194
12.15Experimental method	195
12.16Results	197
12.17Discussion	199
12.18Final remarks	201
12.19Flicker and fibers	201
13 AM Noise	205
13.1 Basics	206
13.2 Single channel measurement	207
13.3 Dual channel (correlation) measurement	209
13.4 Schottky and tunnel diode power detectors	210
13.5 The double balanced mixer	214
13.6 Power detector noise	215
13.7 Design of the front-end amplifier	217
13.8 The measurement of the power detector noise	225
13.9 AM noise in optical systems	226
13.10AM noise in microwave photonic systems	229
13.11Calibration	231
13.12Examples	234
13.13Final remarks	239
14 Exotic Methods	241
15 Applications	243
15.1 A Taste of Phase Noise and Frequency Stability	243
15.2 Radars and radioastronomy	245

16 The companion of the Enrico's Chart	247
16.1 Introduction	247
16.2 Two-Port Components	251
16.3 Oscillators and the Leeson Effect	251
16.4 The Allan Variance	256
16.5 Other Options for the Two-Sample Variance	256
16.6 Relations Between Phase Noise and Variance	257
16.7 Confidence Intervals	257
16.8 Measurement Techniques	257
Notation, Symbols and Acronyms	285

TO DO - The Siccardi amplifier

Chapter 1

Basic Facts¹

Most of this book is about the clock signal, that is, a periodic signal intended as a frequency reference or as a time reference in metrology and in a variety of applications. The clock signal is inevitably affected by AM and PM noise, which breaks the periodicity in strict sense. Fig. 1.1 shows a sinusoidal clock signal

$$v(t) = V_0[1 + \alpha(t)] \cos[2\pi\nu_0 t + \varphi(t)] \quad \text{clock signal} \quad (1.1)$$

where V_0 is the nominal amplitude, ν_0 the frequency, $\alpha(t)$ is the random fractional amplitude, and $\varphi(t)$ is the random phase. Albeit we refer implicitly to a voltage, (1.1) is general, thus it applies to currents, electric field, magnetic fields, etc. It may be wise to replace the symbol $v(t)$ with the appropriate one. Other waveforms are also a widely used, typically the square wave in digital circuits. In most practical cases we can assume that AM and PM noise is very small, or locally small during e measurement. Equation. (1.1) is a simplification of (1.29) because it describes the measured signal, hiding the difference between parametric noise (true PM noise) and additive noise.

¹A significant part of the material in this chapter comes from E. Rubiola, F. Vernotte, “The Companion of Enrico’s Chart...”




Figure 1.1: Clock signal (1.1), observed with an oscilloscope triggered by an external noise-free reference. The shadowed (pink) arrows indicate where the signal is stationary, unaffected by noise. The elliptical shadowed (cyan) regions emphasize where noise shows up most.

1.1 Physical quantities of interest

We introduce the main quantities used to describe the clock signal (1.1). The interplay of the quantities related to phase and frequency (not to amplitude) is represented in Fig. 1.2, together with the relations between power spectral densities (PSDs).

Symbols and fonts of special quantities

Some quantities, like $x(t)$, $y(t)$ and D defined in this Section, have a special meaning in this book. We use the sans serif font, instead of the regular math font commonly found in the literature, to emphasize this, and to set the regular x , y and D free for general use. The same applies to the coefficients b_i , d_i , h_i and k_i defined later.

Amplitude V_0

Neglecting the small effect of AM noise, V_0 is the peak amplitude.

Fractional amplitude fluctuation $\alpha(t)$

The fractional amplitude fluctuation, dimensionless, is defined by (1.1). The IEEE Standard 1139 introduces the clock signal as $v(t) = V_0 + \epsilon(t) \cos(\dots)$, which defines the amplitude fluctuation $\epsilon(t)$. The fractional amplitude fluctuation is $\alpha(t) = \epsilon(t)/V_0$, which is equivalent to our notation.

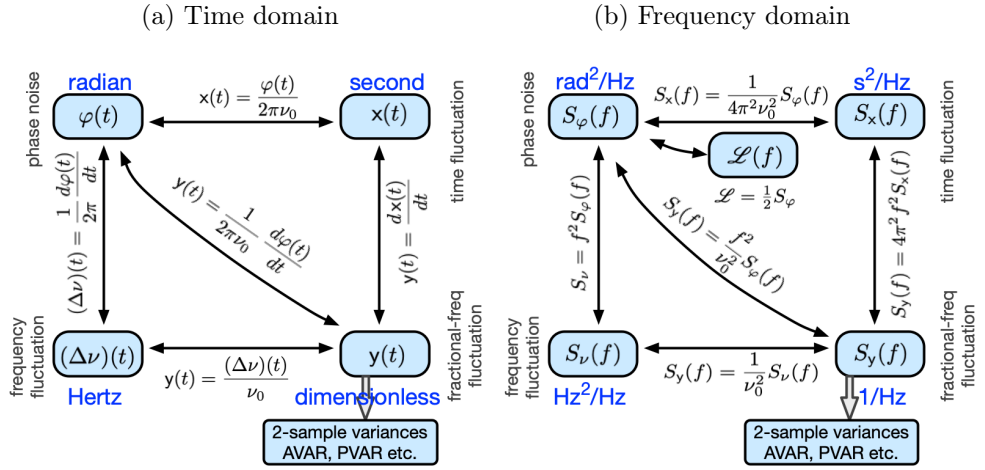


Figure 1.2: Physical quantities which characterize phase and frequency noise (a) in the time domain, and (b) in the frequency domain.

Carrier frequency ν_0

At first reading, we should take ν_0 as the nominal carrier frequency, or its best estimate, so that the randomness goes in $\varphi(t)$. Subtleties will appear later, related to divergent phenomena, and drift.

Fourier (analysis) frequency f

The Fourier frequency is the running variable in spectral analysis, in turn based on the Fourier transform. Unless otherwise specified, f refers to one-sided spectra, where the power associated to negative frequencies is moved to the positive-frequency region.

Angular frequency ω

The angular frequency ω is equivalent to the (regular) frequency, but for the unit rad/s. We use ω as a shorthand for $2\pi\nu$ or $2\pi f$, often without introduced it, with meaning defined by the subscript in the context. For example, ν_b and ω_b may be used interchangeably to denote the frequency, or the angular frequency of a beat note defined by the subscript b . In spectral analysis, ω goes preferably with two-sided Fourier transforms.

Phase fluctuation $\varphi(t)$

The phase fluctuation is the random phase $\varphi(t)$ defined by (1.1), or the equivalent quantity in the case of non-sinusoidal signals. It has the physical dimension of angle, and the appropriate SI² unit is radian, symbol rad. The quantity $\varphi(t)$ results from a variety of physical phenomena impacting on phase and frequency. The PSD of $\varphi(t)$, denoted with

$$S_\varphi(f) \quad [\text{rad}^2/\text{Hz}]$$

²International system.

is the main quantity that should be used to describe the spectrum of phase noise. Its physical dimension is a square angle per unit of frequency, unit rad^2/Hz .

Phase time fluctuation $x(t)$

The phase time fluctuation, or for short phase time, is the phase fluctuation converted into time

$$\text{definition: } x(t) = \frac{\varphi(t)}{2\pi\nu_0} \quad [\text{s}] \quad (1.2)$$

Its PSD $S_x(f)$ follows from the definition of $x(t)$ using the property of the Fourier transform that the derivative operator d/dt maps into a multiplication by $j2\pi f$, where $j^2 = -1$, thus into a multiplication by $4\pi^2 f^2$ in the PSD because of (1.17). Accordingly,

$$S_x(f) = \frac{1}{4\pi^2\nu_0^2} S_\varphi(f) \quad [\text{s}^2/\text{Hz} \equiv \text{s}^3]. \quad (1.3)$$

Frequency fluctuation $(\Delta\nu)(t)$

Because frequency is the derivative of the phase, we can express $\varphi(t)$ as a frequency fluctuation

$$\text{definition: } (\Delta\nu)(t) = \frac{1}{2\pi} \frac{d\varphi(t)}{dt} \quad [\text{Hz}] \quad (1.4)$$

which leads to the alternate representation of (1.1)

$$v(t) = V_0 [1 + \alpha(t)] \cos[2\pi\nu_0 t + \int (\delta\nu)(t) dt]. \quad (1.5)$$

The notation $(\Delta\nu)(t)$, with $\Delta\nu$ in parentheses, emphasizes the fact that $(\Delta\nu)$ is an unbreakable quantity, function of time.

The PSD is found using the property that $d/dt \rightarrow \times 4\pi^2 f^2$, thus

$$S_\nu(f) = f^2 S_\varphi(f) \quad [\text{Hz}^2/\text{Hz} \equiv \text{Hz}]. \quad (1.6)$$

Albeit the notation $S_{\Delta\nu}$ is more common than S_ν , the latter should be preferred because the PSD is already insensitive to the constant ν_0 , by definition, thus it shows only the fluctuation $\nu - \nu_0$.

Fractional frequency fluctuation $y(t)$

The fractional frequency fluctuation is

$$\text{definition: } y(t) = \frac{(\Delta\nu)(t)}{\nu_0} \quad [\text{dimensionless}] \quad (1.7)$$

It follows from (1.2), (1.4) and (1.7) that

$$y(t) = \frac{dx(t)}{dt} \quad (1.8)$$

and

$$y(t) = \frac{1}{2\pi\nu_0} \frac{d\varphi(t)}{dt}. \quad (1.9)$$

The PSD is found using $d/dt \rightarrow \times 4\pi^2 f^2$, thus

$$S_y(f) = \frac{f^2}{\nu_0^2} S_\varphi(f) \quad [\text{Hz}^{-1} \equiv \text{s}]. \quad (1.10)$$

Fractional frequency drift D

The fractional frequency drift is defined as

$$\text{definition:} \quad D = \frac{dy(t)}{\nu dt} \quad [\text{s}^{-1} \equiv \text{Hz}], \quad (1.11)$$

assuming a linear drift over time. This choice originates from the fact that aging in most oscillators shows up as a linear drift, or as a drift which is nearly linear for long time, like years. Besides, the drift makes sense only if it is averaged on a suitably long time interval, so that it exceeds the random phenomena. For this reason, the drift is often expressed using unit of time longer than the second, typically day^{-1} or year^{-1} . For example, the aging found in the specs of a 10 MHz OCXO is $2 \times 10^{-10}/\text{day}$.

1.1.1 Extended representation of phase and frequency noise

Working with digital systems, we are regularly faced with phase exceeding $\pm\pi$ because IQ detectors and digital dividers keep record of multiple cycles of the carrier. For this purpose, we find it useful to describe the clock signal with the quantity written in boldface, which is the sum of the deterministic (or nominal) quantity plus the fluctuation

$$\text{phase:} \quad \boldsymbol{\varphi}(t) = 2\pi\nu_0 t + \varphi(t) \quad (1.12)$$

$$\text{frequency:} \quad \boldsymbol{\nu}(t) = \nu_0 + (\Delta\nu)(t) \quad (1.13)$$

$$\text{time:} \quad \boldsymbol{x}(t) = t + \mathbf{x}(t) \quad (1.14)$$

$$\text{fractional frequency:} \quad \boldsymbol{y}(t) = 1 + \mathbf{y}(t). \quad (1.15)$$

The quantity $\mathbf{x}(t)$ is the most obvious. To the layman, $\mathbf{x}(t)$ is the readout of a clock, which is the sum of the ‘exact time’ t plus the ‘error’ $\mathbf{x}(t)$. A true layman would not consider relativity here, and would have no idea about the technical meaning of words like ‘error’ and ‘uncertainty.’ The quantity $\boldsymbol{\nu}(t)$ is the instantaneous frequency, measured in a sufficiently short time, $\boldsymbol{\varphi}(t)$ is the total phase accumulated after IQ detection, and $\mathbf{y}(t)$ differs from 1 by the small fractional fluctuation $\mathbf{y}(t)$.

1.1.2 Average of $\alpha(t)$ and $\varphi(t)$

Most readers may be led to believe that $\alpha(t)$ and $\varphi(t)$ have zero average, or to take these conditions as necessary, because introductory education in statistics is generally restricted to zero average processes in order to escape from the difficulty of divergence. Similarly, the reader may be led to take V_0 and ν_0 as the ‘true value’ of amplitude and frequency. Let us reconsider these hypotheses in a simple way.

In systems of practical interest we observe quite small amplitude and frequency fluctuations, thus we can assume that $|\alpha| \ll 1$. We have measured $|\alpha|$ of the order




Figure 1.3: Physical concept of power spectral density.

of $10^{-7} \dots 10^{-4}$ in quartz oscillators and in synthesizers [Rubiola and Boudot, 2007], and $\langle \alpha(t) \rangle$ averaged in the full measurement time is of the same order.

The hypothesis that $|\varphi(t)| \ll 1$ rad makes sense only for two-port systems, where $\varphi(t)/2\pi\nu_0$ is the fluctuation of the input-to-output delay. The same hypothesis does not fit the behavior of oscillators, which always contains divergent processes. Even the most sophisticated quantum oscillators accumulate multiple-period timing errors if we allow sufficient time to do so.

Measuring oscillators, we often rely on analog phase meters, which require that $|\varphi(t)| \ll 1$ rad for proper operation, and we measure the error signal after phase-locking a reference oscillator to the oscillator under test. If the phase lock functions normally, the average of the error signal is generally close to zero. Otherwise, we may use digital phase meters, which rely on the in-phase and quadrature (I-Q) of the clock signal versus a reference. I-Q detection enables the measurement of large phase angles, not limited to $\pm\pi$. In both cases the measurement of large phase angles is possible, but we generally need that $\dot{\varphi}(t)$ results in a small angular frequency, so that $|\dot{\varphi}| \ll 2\pi\nu_0$. Under this hypothesis, ν_0 is a good approximation of the carrier frequency for any practical purpose. The value of $|\dot{\varphi}|/2\pi\nu_0$ spans from 10^{-4} for cheap watches to 10^{-16} (accuracy) and 10^{-18} (fluctuations) for the frequency-standard prototypes found in metrology labs.

1.1.3 Bandwidth of $\alpha(t)$ and $\varphi(t)$

The bandwidth B of $\varphi(t)$ and $\alpha(t)$ deserves some attention. Because modulation needs sidebands with appropriate symmetry around ν_0 , the theoretical maximum B is equal to ν_0 . That said, in numerous cases of interest the quantity B/ν_0 is rather small, likely of $10^{-4} \dots 10^{-2}$, and not necessarily the same for $\varphi(t)$ and $\alpha(t)$. The autocorrelation function of noise spans over a time of the order of $1/B$, or equivalently ν_0/B oscillations. Thus, cycle-to-cycle noise is too small to be visible. This is why we use the external trigger in Fig. 1.1. Otherwise, triggering on signal itself would require a dual time base with an ideally stable delay $> 1/B$.

1.2 Power spectral density (PSD)

The *variance*³ (generalized power) of a quantity q , denoted by σ^2 , is the mean square of q . Formally, $\sigma^2 = \mathbb{E}\{|q - \mu|^2\}$, where $\mathbb{E}\{\}$ is the mathematical expectation, and $\mu = \mathbb{E}\{q\}$ is the average. The *power spectral density* (PSD), denoted with $S(f)$, tells us how σ^2 is distributed in frequency, as illustrated in Fig. 3.1. The variable f is called ‘Fourier frequency’ to differentiate it from the carrier frequency (constant). The single-sided PSD ($f > 0$) is generally preferred to the two-sided PSD with no need of saying so. Referring to the quantity q , the subscript q is optionally added, as in σ_q^2 and $S_q(f)$.

In practical measurements the PSD $S_q(f)$ of the random function $q(t)$ is best *evaluated* as

$$\widehat{S}_q(f) = \frac{2}{T} \left\langle Q_T(f) Q_T^*(f) \right\rangle_m, \quad (1.16)$$

where the hat accent means ‘estimator’ of S_q (the reader may ignore it at first reading), $Q_T(f)$ is the Fast Fourier Transform of $q(t)$ sampled and truncated on an appropriate duration T , the superscript ‘*’ means complex conjugate, the $\langle \rangle_m$ operator is the average on m acquisitions, and the factor of 2 is needed for energy conservation after deleting the negative frequencies⁴. Equation (1.16) is used in the Welch algorithm for the estimation of power spectra [Welch, 1967]. The optional data overlapping used in the Welch algorithm, and the optional window function are not explicit in (1.17). The most popular window functions are known under the names Bartlett, Blackman-Harris, flat-top, Hamming, Hann, Parzen and of course Welch. The quantity $S(f) = \dots$ defined by (1.16), with optional windowing but with no hat and no averaging ($m = 1$), is called *periodogram*.

Measuring phase noise, $S_\varphi(f)$ is the quantity that should be used. Thus, following (1.16), $S_\varphi(f)$ is best *evaluated* as

$$\widehat{S}_\varphi(f) = \frac{2}{T} \left\langle \Phi_T(f) \Phi_T^*(f) \right\rangle_m \quad [\text{rad}^2/\text{Hz}], \quad (1.17)$$

1.2.1 Deeper thoughts about the PSD

In statistics, the PSD $S(f)$ is formally defined as the Fourier transform of the autocovariance of a *random process*. In turn, the random process is a set of sample functions or distributions, called *realizations*, each of which is indexed by one outcome of a random experiment. Time statistics and ensemble statistics are different concepts, and they are interchangeable only in the case of ergodic processes.

In contrast, some Authors take (1.17) as the definition of $S_\varphi(f)$, with no ‘hat.’ The problem with this choice is that it avoids some key concepts of statistics, making the uncertainty difficult to understand.

³The two sample variance we find later is a specialized instance of this general concept.

⁴Because $q(t)$ is a real function, its Fourier transform $Q(f)$ is Hermitian function, that is, $Q(f) = Q^*(-f)$. Accordingly, $\Re\{Q(f)\}$ is even function of f , and $\Im\{Q(f)\}$ is odd function of f . Thus, all the information is contained in the $f > 0$ half-plane, and the $f < 0$ half-plane is redundant and can be deleted. In the case of the FFT, such redundant region is generally mapped to $f_s/2 > f > f_s$, where f_s is the sampling frequency.

Using the mathematical concept of process requires one to define the random experiment. Since mathematics does not provide general rules for this, we need to make a choice. For example, we may identify the random experiment with a large abstract class, or more pragmatically with the action of picking one oscillator from preproduction samples, from a batch, or from a larger part of a production. In turn, a realization may be identified with the waveform $\varphi(t)$ obtained by comparing such oscillator with a noise-free reference. This opens deeper questions, like the meaning and the scientific legitimacy of ‘typical’ spectra (also stability, aging, drift, and other parameters) found in data sheets.

As a matter of fact, the right-hand side of (1.17) can always be calculated from experimental data. The question arises, whether or not the estimator converges to the PSD

$$\widehat{S}_\varphi(f) \rightarrow S_\varphi(f) \quad \text{for large } m. \quad (1.18)$$

This is true for stationary processes (the statistical properties are independent of the origin of time). That said, relevant processes often found in oscillators, like flicker and random walk of frequency, are not stationary in a strict sense. Fortunately, evaluating (1.17) such processes can be treated as stationary.

Finally, we notice some analogies between *stationarity* and *repeatability*, and also between *ergodicity* and *reproducibility*, with the caveat that stationarity and ergodicity are mathematical concepts, while repeatability and reproducibility are defined by the International Vocabulary of Metrology VIM [Joint Committee for Guides in Metrology (JCGM), 2012, Entry 2.21 and 2.25] and have quite a technical meaning related to experimental outcomes.

1.2.2 The polynomial law, or power law

The polynomial fit, known as the *power law* or *polynomial law*, is widely used to model phase noise and related quantities. It is often written as

$$S_\varphi(f) = \sum_{n=-4}^0 \mathbf{b}_n f^n, \quad (1.19)$$

$$S_x(f) = \sum_{n=-4}^0 \mathbf{k}_n f^n, \quad (1.20)$$

where the values of n correspond to the noise types listed in Table 1.1. These noise types are found in oscillators, with additional negative-exponent terms sometimes needed, $n < -4$. Limitations apply to two-port devices because the input-to-output delay is not allowed to diverge (CF Sec. 16.2 and 16.3).

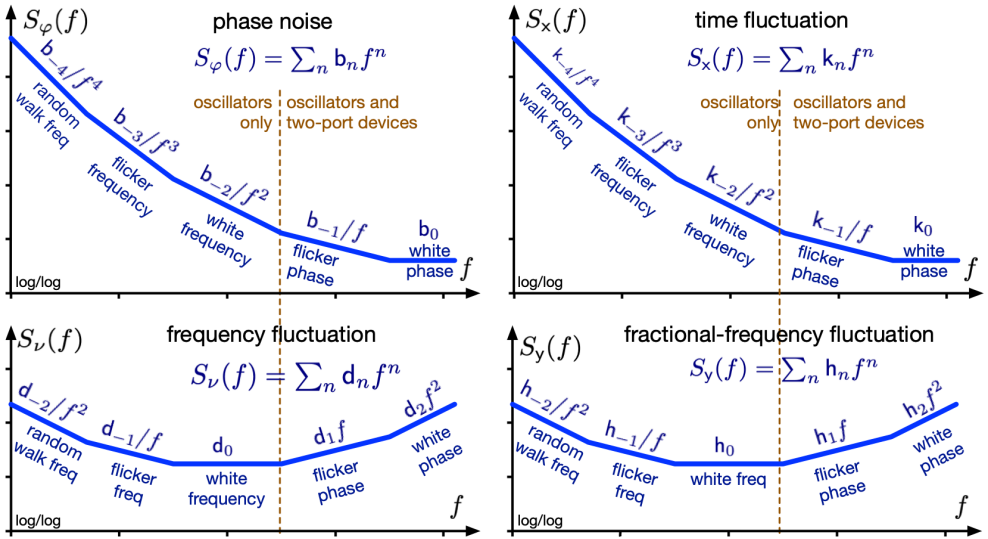


Figure 1.4: .

Transposing (1.19)-(1.20) to frequency noise, the polynomial law is written as

$$S_\nu(f) = \sum_{n=-2}^2 d_n f^n, \quad (1.21)$$

$$S_y(f) = \sum_{n=-2}^2 h_n f^n. \quad (1.22)$$

Notice that, for a given process the exponents of f differ by 2 from (1.19)-(1.20) to (1.21)-(1.22), in agreement with the bounds of the sum.

In proper mathematical terms, (1.19)-(1.22) are Laurent polynomials, which is the extension of the regular polynomials to negative powers of the variable.

1.3 The deprecated quantity $\mathcal{L}(f)$

Most often in the technical literature and data sheets, the term *phase noise* refers to the quantity $\mathcal{L}(f)$, defined since the IEEE Standard 1139-1988 [Hellwig, 1988] as

$$\text{definition: } \mathcal{L}(f) = \frac{1}{2} S_\varphi(f). \quad (1.23)$$

Plots and numerical values are always given as

$$10 \log_{10}(\mathcal{L}(f)) \quad [\text{dBc/Hz}]. \quad (1.24)$$

Table 1.1: Polynomial law and basic types of noise

Noise type	$S_\varphi(f)$	$S_x(f)$	$S_\nu(f)$	$S_y(f)$
Blue phase noise*	b_1	k_1	$d_3 f^3$	$h_3 f^3$
White phase noise	b_0	k_0	$d_2 f^2$	$h_2 f^2$
Flicker phase noise	$\frac{b_{-1}}{f}$	$\frac{k_{-1}}{f}$	$d_1 f$	$h_1 f$
White frequency noise	$\frac{b_{-2}}{f^2}$	$\frac{k_{-2}}{f^2}$	d_0	h_0
Flicker frequency noise	$\frac{b_{-3}}{f^3}$	$\frac{k_{-3}}{f^3}$	$\frac{d_{-1}}{f}$	$\frac{h_{-1}}{f}$
Frequency random walk	$\frac{b_{-4}}{f^4}$	$\frac{k_{-4}}{f^4}$	$\frac{d_{-2}}{f^2}$	$\frac{h_{-2}}{f^2}$
Integrated flicker FM noise*	$\frac{b_{-5}}{f^5}$	$\frac{k_{-5}}{f^5}$	$\frac{d_{-1}}{f^3}$	$\frac{h_{-1}}{f^3}$
Integrated RW frequency*	$\frac{b_{-6}}{f^6}$	$\frac{k_{-6}}{f^6}$	$\frac{d_{-2}}{f^4}$	$\frac{h_{-2}}{f^4}$

★ Process not shown in Fig. 1.4

After over 30 years in the field, the author believes that $\mathcal{L}(f)$ is misleading and, if the community started from scratch, $S_\varphi(f)$ would be used instead. The roots are found in a one-time symposium co-organized by IEEE and NASA in 1964 [Chi, 1964].

Because $\varphi(t)$ is an angle, the dimension of $S_\varphi(f)$ is time⁵. Accordingly, the appropriate unit is rad for $\varphi(t)$ and rad²s for $S_\varphi(f)$, but the equivalent unit rad²/Hz is generally preferred. In logarithmic units⁶ we use

$$10 \log_{10} S_\varphi(f) \quad [\text{dBrad}^2/\text{Hz}]. \quad (1.25)$$

It follows from (1.23) that $S_\varphi(f)$ and $\mathcal{L}(f)$ have the same physical dimension⁷

$$\dim(\mathcal{L}(f)) = \dim(S_\varphi(f)) = \mathsf{T} \quad (1.26)$$

but different units, like a mass in kg or in lb. It is worth objecting that there is no reason to change symbol after switching unit, like in $M = 1.5$ kg and $M = 3.3$ lb. So,

⁵According to the SI Brochure (Sec. 2.3.3), angles, when expressed in radians and steradians respectively, are treated as quantities with the unit one (no unit). Nonetheless, the symbols rad is written explicitly where appropriate, in order to emphasize that the quantity being considered is an angle. **A paragraph about the unit of angle should be added, somewhere else.**

⁶The decibel, symbol dB, is a non-SI unit accepted for use with the SI.

⁷Note that the angle has a unit, the radian, symbol rad, but has no physical dimension.

why should we change the symbol for phase noise from S_φ to \mathcal{L} just because of the unit of angle? Anyway, the unit associated to $\mathcal{L}(f)$ is \mathfrak{A}^2/Hz where \mathfrak{A} is the fictitious unit of angle⁸ defined as

$$\mathfrak{A} = \sqrt{2} \text{ rad} \simeq 81^\circ. \quad (1.27)$$

The unit dBc/Hz is even more confusing. It is obvious that ‘c’ cannot be read ‘referred to the carrier,’ as most people have in mind. Instead, taking (1.23)-(1.24) literally, ‘c’ is a square unit of angle,

$$\text{c} \equiv \mathfrak{A}^2 = 2 \text{ rad}^2.$$

Notwithstanding this, ‘c’ and ‘c/Hz’ alone, with no ‘dB,’ are never seen in the literature, and no unit of angle associated with $\mathcal{L}(f)$ either.

We advocate abandoning $\mathcal{L}(f)$ because $\mathcal{L}(f)$ is a non SI unit. There are non-SI units accepted for use with the SI because they are useful or ubiquitous, listed in Table 8 of the SI brochure (day, litre, electronvolt, etc.). In contrast, the unit dBc/Hz is even not mentioned among such ‘tolerated’ non-SI units, thus it falls in the category of the *deprecated units* described in SI brochure at the end of Section 4, verbatim

There are many more non-SI units, which are either of historical interest, or are still used in specific fields (for example, the barrel of oil) or in particular countries (the inch, foot and yard). The CIPM can see no case for continuing to use these units in modern scientific and technical work.

The quantity $S_\varphi(f)$ should be used instead, which is consistent with the International System of Units SI [Bureau International des Poids et Mesures (BIPM), 2019]—with the minor caveat that the decibel is a non-SI unit accepted for use with the SI units. In order to prevent discontinuity in notation and units, $\mathcal{L}(f)$ may still be kept as an auxiliary scale.

1.3.1 The former, obsolete definition of $\mathcal{L}(f)$

In the early time, $\mathcal{L}(f)$ was defined in terms of SSB noise (Fig. 1.5) as

$$\mathcal{L}(f) = \frac{\text{noise power in 1 Hz bandwidth}}{\text{carrier power}} \quad (\text{wrong!}), \quad (1.28)$$

and always given as $10 \log_{10} \mathcal{L}(f)$ in dBc/Hz. The historical meaning of the symbol ‘c’ is ‘referred to the carrier.’ For example, -120 dBc/Hz means that ‘the noise sideband in 1 Hz bandwidth is 120 dB below the carrier power.’ Sadly, more than 30 years after the first version of the IEEE Standard 1139 [Hellwig, 1988], the old definition is still in the mind of numerous engineers and physicists. A few problems with (1.28) have been identified, discussed below.

⁸There is no such a unit of angle, and the choice of a nearly unreadable symbol for it is intentional.




Figure 1.5: The obsolete and incorrect definition of $\mathcal{L}(f)$.

The obsolete definition of $\mathcal{L}(f)$ is a conceptually incorrect definition of PM noise

The power spectrum does not say if the sidebands relate to AM noise, to PM noise, or to any combination of. This concept illustrated in Fig. 1.6, which shows the simple case of a carrier affected by white band-limited noise

$$(A) \quad v(t) = V_0 \cos[\omega_0 t + \varphi(t)] \quad \text{PM noise}$$

$$(B) \quad v(t) = V_0 [1 + \alpha(t)] \cos(\omega_0 t) \quad \text{AM noise}$$

Looking at the Fourier transforms $V(\omega)$, the difference is that in PM (A) the sidebands have *odd* symmetry around ω_0 in the real part and *even* symmetry in the imaginary part, while in AM (B) the sidebands have *even* symmetry in the real part and *odd* symmetry in the imaginary part. However, the power spectra are identical

$$S_v(f) \propto V_0^2 [\Re^2\{V\} + \Im^2\{V\}] , \quad f = \omega/2\pi, \quad f > 0 ,$$

thus there no way to identify the modulation type from the spectrum.

That said, it is often possible to assess that the noise sidebands of Fig. 1.5 are PM noise on the ground of physical reasoning. For example, $1/f^3$ and $1/f^4$ processes are present in the PM noise of oscillators, not in AM noise. Similarly, there is a strong dominance of PM noise at the output of high-order multiplication, with negligible contribution of AM noise. However, inferring PM noise from a spectrum that contains both AM and PM is a poor replacement for a true measure of PM noise.

The term ‘offset frequency’ should be replaced with ‘Fourier frequency’

The running variable f in $S_\varphi(f)$ and $\mathcal{L}(f)$ is often called ‘offset frequency.’ Such term is inappropriate and misleading because it refers to one sideband, while two sidebands with the appropriate symmetry are necessary to form a modulated signal. Instead, f should be called ‘Fourier frequency’ in proper mathematical language, or ‘modulation frequency’ in engineering jargon.

The obsolete definition of $\mathcal{L}(f)$ is experimentally incorrect

All phase noise analyzers use a phase detector to assess the random phase $\varphi(t)$ before evaluating the PSD. So, they truly measure $S_\varphi(f)$, not $\mathcal{L}(f)$, and derive $\mathcal{L}(f)$ using $\mathcal{L}(f) = \frac{1}{2} S_\varphi(f)$.

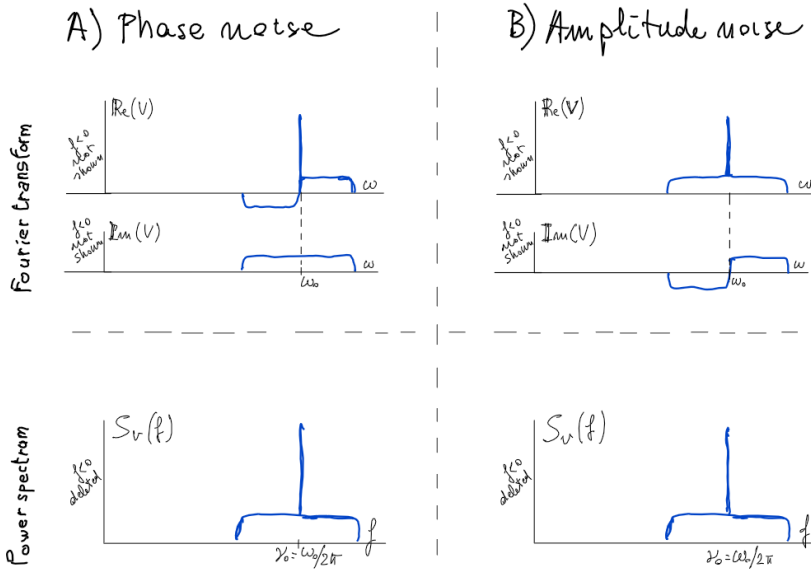


Figure 1.6

Indeed, the confusion between the true measurement of PM noise and its description in terms of SSB-noise like in Fig. 1.5 was already present in the 1964 IEEE NASA Symposium⁹. One attempt was made to measure the SSB noise¹⁰, using a complex experimental method which requires to suppress the carrier by adding an equal and opposite reference signal phase-locked to the input. It is only after suppressing the carrier that the signal fits in the limited dynamic range of a regular spectrum analyzer. However, this method seems to have disappeared from the literature and from the laboratory practice.

We can only infer a reason for the confusion mentioned, that the mathematical theory of random angular modulations was not sufficiently known in the community of experimentalists.

The obsolete definition of $\mathcal{L}(f)$ is unsuitable to large angles

Large phase swings are always present in oscillators at low Fourier frequencies (long measurement time) and in optics, where the carrier frequency is high. When $\int S_\varphi(f) df$ integrated over the full spectrum approaches or exceeds $1 \text{ rad}^2/\text{Hz}$, (1.28) breaks down because the noise sidebands in angular modulations come at the expense of the carrier power. Conversely, there is no reason to question the validity of (1.17) and (1.23), even in the case of huge number of cycles. In the presence of large angles, the correct measurement of $S_\varphi(f)$ is only a matter of hardware design.

⁹citation

¹⁰Citation. Horn?

The ‘dBc/Hz’ is a deprecated unit of phase noise

The unit dBc/Hz is deprecated because (i) it does not describe a phase modulation, deterministic or random, and (ii) if used to describe phase noise, is incompatible with the SI.

1.3.2 A final word

The power ratio as defined in the right-hand side of (1.28) is definitely not evil. In some domains the noise per unit of bandwidth matters, regardless of the physical origin. This is the case of the reciprocal mixing in the superheterodyne receiver [Razavi, 2012, Sec. 8.7.2], where a strong emission in neighbor channels leaks into the IF because of the local-oscillator noise sidebands. Similarly, Doppler radar visibility of subclutter targets is typically determined by noise in the oscillator’s sideband power density. See [Leeson and Johnson, 1964, Leeson and Johnson, 1966, Grebenkemper, 1981], and [Barton and Leonov, 1998, Skolnik, 2008] as general references on radars. As noted previously, due to amplitude limiting, sideband noise in the cited examples is generally dominated by phase noise, leading to the persistence of the spectrum sideband usage in these applications.

However, a problem of jargon is to be solved. Albeit (1.28) is a fully legitimate measure of the SSB noise ratio, correctly measured as a dimensionless ratio or expressed in dBc/Hz, everybody takes the quantity $\mathcal{L}(f)$ and the unit dBc/Hz for phase noise.

1.4 Two-port components

1.4.1 Additive and parametric noise

People tend to find these concepts either quite simple, or rather confusing. The point is that (1.1) describes the clock signal as it is observed, hiding the physics of noise.

To understand, we start with the example of a noise-free radio broadcasting $v_e(t) = V_e[1 + \alpha_m(t)] \cos[2\pi\nu_0 t + \varphi_m(t)]$, where $\alpha_m(t)$ is AM, and $\varphi_m(t)$ is PM. AM and PM may be present simultaneously, as in the old analog television¹¹ and in QAM modulations. The *received* signal

$$v(t) = V_0[1 + \alpha_m(t)] \cos[2\pi\nu_0 t + \varphi_m(t)] + n(t) \quad (1.29)$$

includes the noise $n(t)$, which is the receiver’s own noise, atmospheric noise, and other forms of noise collected by the antenna. However, the signal is *detected* as

$$v(t) = V_0[1 + \alpha(t)] \cos[2\pi\nu_0 t + \varphi(t)]. \quad (\text{same as (1.1)})$$

In fact, a sinusoidal signal at a given frequency has only two degrees of freedom, amplitude and phase. Thus, $\varphi(t)$ is the sum of $\varphi_m(t)$ plus the image of $n(t)$ after

¹¹More precisely, TV audio is FM for compatibility with audio broadcasting, which is equivalent to PM.




Figure 1.7: Additive (SSB) noise and parametric (AM and PM) noise. Bob Nelson, Keysight.

phase detection. Similarly, $\alpha(t)$ is the sum of $\alpha_m(t)$ plus the image of $n(t)$ after amplitude detection. Random $n(t)$ is called *additive noise*, while random $\alpha_m(t)$ and $\varphi_m(t)$ are called *parametric noise*. The main differences are

- Parametric noise originates from the low-frequency signals $\varphi_m(t)$ and $\alpha_m(t)$ up converted to sidebands, while additive noise originates around ν_0 .
- The power of $n(t)$ is equally split into AM and PM noise. In contrast, there is no a priori reason for $\alpha_m(t)$ and $\varphi_m(t)$ to give the same contribution to the spectrum.

If $\varphi_m(t)$ and $\alpha_m(t)$ are small¹² and have bandwidth $B_m \ll \nu_0$ their image in the RF spectrum is a pedestal $\pm B_m$ wide centered at ν_0 .

It is worth mentioning that white AM and PM noise results from white $n(t)$, and from white $\alpha_m(t)$ and $\varphi_m(t)$. Therefore, the common belief that white AM/PM noise is only additive is incorrect. Figure 1.8 provides an example of oscillator where the white-noise region contains both additive and parametric noise. Conversely, ‘colored’ noise ($1/f$, $1/f^2$, etc.) is generally of parametric origin. A narrow-band $n(t)$ well centered at ν_0 appearing as $1/f$ noise is totally unrealistic, albeit conceptually possible.

1.4.2 The unfortunate term ‘added noise’

The term ‘added noise’ is seen in commercial phase noise analyzers to denote the phase noise *added* by a two-port component under test, usually an amplifier, when the component is inserted in the signal path. This is an unfortunate choice because

¹²More precisely, small modulation angle $|\varphi_m(t)| \ll 1$ is necessary for higher-order sidebands to be negligible, and the PM to be accurately described by the first USB and LSB. Conversely, only $|\alpha_m(t)| < 1$ is formally necessary, but the $|\alpha_m(t)| \ll 1$ condition makes sense in real systems.




Figure 1.8: AM and PM noise of a microwave oscillator. Looking at the white-noise region ($f \gtrsim 10$ kHz), we see a significant difference between AM noise and PM noise. This is only possible if the white region contains a certain amount of parametric noise. Courtesy of Ulrich L. Rohde.

the term ‘added’ is easily mistaken for ‘additive.’ Of course, the ‘added’ noise consists of white and flicker PM noise, thermal drift, aging, etc., while the ‘additive’ noise is white or quasi white.

In proper terminology, added noise refers to a *differential measurement* because the instrument measures the quantity $\varphi - \psi$, i.e., output minus input, and its statistical properties (spectra and variances, at will). The oscillator noise ψ is a common-mode signal, thus it is rejected.

1.4.3 White noise

Given a white noise $n(t)$ of PSD N [W/Hz] added to a carrier of power P , the white PM noise’s PSD is

$$b_0 = \frac{N}{P} \quad [\text{rad}^2/\text{Hz}]. \quad (1.30)$$

Radio engineers often express the noise in terms of noise factor F . Accordingly, N is expanded as $N = FkT_0$, where $kT_0 = 4 \times 10^{-21}$ W/Hz is the thermal energy at the standard temperature $T_0 = 290$ K (17 °C). The Noise Figure is defined by $\text{NF} = 10 \log_{10}(F)$. For reference, the phase noise of a noise-free device ($F = 1$, or $\text{NF} = 0$ dB) in the presence of a 1-mW carrier (0 dBm) is 4×10^{-18} rad²/Hz, thus -174 dB rad²/Hz or -177 dBc/Hz. Besides thermal energy, $n(t)$ may originate from

shot noise. In our experience, experimentalists in optics are led to confusion by the fact that because noise is represented as a temperature, while in optics the thermal energy is generally negligible.

Additional white PM noise, not accounted for in (1.30) can result from parametric effects, as explained in Sec. 16.2

1.4.4 Flicker Noise

Flicker noise has PSD proportional to $1/f$. It has been observed that the flicker PM of RF and microwave amplifiers is rather constant vs P and ν_0

$$b_{-1} = C \quad (\text{constant vs } P \text{ and } \nu_0). \quad (1.31)$$

It follows from (1.30) and (1.31) that the corner frequency $f_c = b_{-1}/b_0$, where flicker equals white noise, depends on P .

Albeit the common integral $\int_a^b (1/f) df = \ln(b/a)$ diverges for $a \rightarrow 0$ or $b \rightarrow \infty$, the practical result is surprisingly small. To convince the reader, we evaluate $\ln(b/a)$ for the largest conceivable bounds, from the reciprocal of the age of the universe ($a = 2.3 \times 10^{-18}$ Hz) to the reciprocal of the Planck time ($b = 1.9 \times 10^{43}$ Hz). The results is $\ln(b/a) = 140.3$, i.e., 21.5 dB. So, if the flicker coefficient is $k_{-1} = 10^{-24} \text{ s}^2$, ($\sqrt{k_{-1}} = 1 \text{ ps}$), the total $1\text{-}\sigma$ fluctuation is $\sqrt{140.3 \times (10^{-24} \text{ s}^2)} = 11.8 \text{ ps}$.

The general literature suggests that the spectrum of flicker is $1/f^\eta$ with η close to one. It turns out that the input-to-output delay never grows too large even for $\eta > 1$. Try this yourself with $\eta = 1.1$ and the integration bounds from 10^{-9} Hz (the reciprocal of 30 years, the supposed device's lifetime) to 10^9 Hz noise bandwidth (above the highest Fourier frequency found in any commercial noise analyzer).

1.4.5 Input-to-Output Delay

The fluctuation ($1\text{-}\sigma$) of the input-to-output delay of a two-port device is given by

$$\delta T = \sqrt{\int_a^b S_x(f) df}, \quad (1.32)$$

which follows from the time-fluctuation PSD integrated on the appropriate bandwidth $[a, b]$. Common sense suggests that δT does not diverge, nor does it grow disproportionately during the device life. This limits the phase noise to finite-bandwidth white PM noise and to flicker PM noise.

Environmental parameters, like humidity and thermal drift, have only localized effect in time, or are periodic. Random walk, aging and other ever growing phenomena are possible, but their amount is generally small enough not to significantly affect the delay over the lifetime of the device. Electrical engineers may be familiar with similar effects in voltage references or in the offset of analog components.

1.5 Scaling the carrier frequency up and down




Figure 1.9: Example of amplifier PN

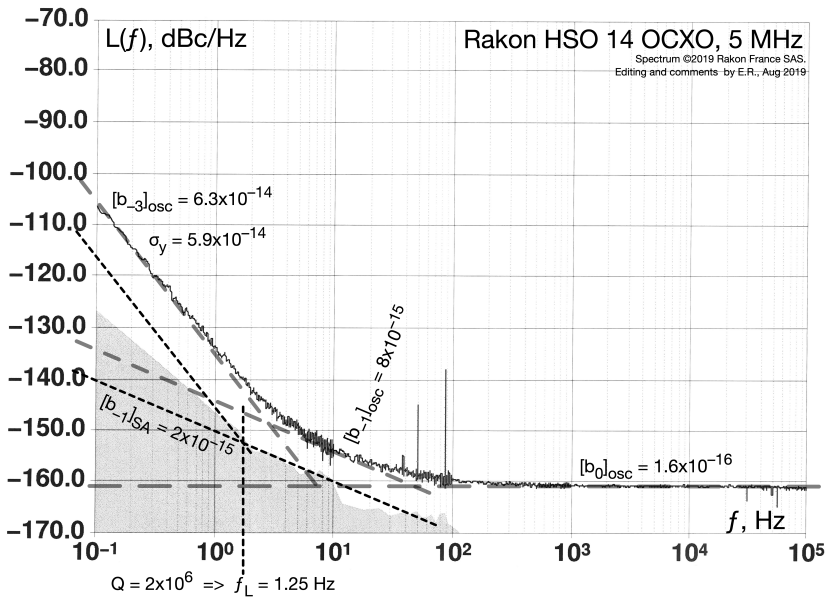


Figure 1.10: Example of oscillator PN

1.5.1 The quantities φ , x , $\Delta\nu$ and y in frequency synthesis

The ideal, noise-free synthesizer is the electrical analogue of a play-free gearbox. It delivers an output frequency $\nu_o = (\mathcal{N}/\mathcal{D})\nu_r$, where \mathcal{N}/\mathcal{D} is the rational number which defines the synthesis ratio, and ν_r is the reference frequency. Thus, the synthesizer transfers the quantities $x(t)$ and $y(t)$ from the reference input to the output, un-




Figure 1.11: Phase noise and phase time fluctuations in frequency synthesis.

changed. For example, shifting the reference by +1.2 ppm, the output frequency will be +1.2 ppm off the nominal value, thus 150 Hz higher if the output is set to 125 MHz. Similarly, introducing a 100 ps delay with a line stretcher at the input results in the output shifted by 100 ps. This value is the same at 5 MHz and 125 MHz output frequency.

By contrast, the synthesis is ruled by $\varphi = (N/D)\psi$ when we express the phase shift as an angle, ψ at the input and φ at the output. For example, a +1 mrad shift of the 10 MHz reference results into a +12.5 mrad shift if the output is set to 125 MHz. Accordingly, the phase noise spectrum is ruled by

$$S_{\varphi}(f) = \left(\frac{N}{D}\right)^2 S_{\psi}(f). \quad (1.33)$$

The above statements are simplistic, to the extent that we have not included the dynamic behavior, the noise bandwidth, and other phenomena. The following limitations apply.

Digital dividers

In digital signals, phase noise exists only on rising and falling edges, thus it is sampled at $2\nu_0$, twice its own frequency. Frequency division $\div D$ results in lower sampling frequency, which originates aliasing. Aliasing impacts only white noise because folding multiple aliases proportional to $1/f$ has negligible effect. Thus, the white PM noise of the divided signal is ruled by $S_{\varphi}(f) = (1/D)S_{\psi}(f)$ instead of $(1/D^2)S_{\psi}(f)$.

Output stage

For $N/D \ll 1$, the scaled-down phase noise may hit the phase noise of the output stage. When this happens, $S_{\varphi}(f)$ is limited by the output stage.

High-order multiplication

Angular modulations are ruled by the property that the total power is constant, thus the sideband power comes at expense of the carrier power. Carrier and n -th sideband amplitudes are described by the Bessel functions $J_n(m)$. For $N/D \gg 1$, the random

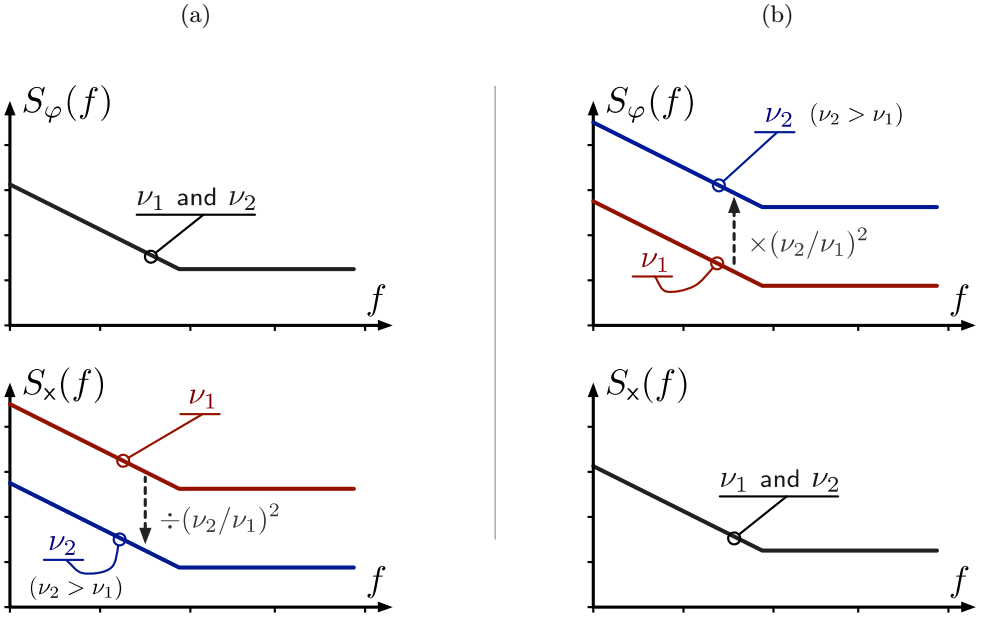


Figure 1.12: Basic phase-noise behavior of components and systems, (a) Phase-type, and (b) Time-type. We show how the carrier frequency affects, or does not affect $S_\varphi(f)$ and $S_x(f)$. The effect of aliasing, not shown here, may change the white-noise scaling from $(\nu_2/\nu_1)^2$ into ν_2/ν_1 . In contrast, flicker and other colored noise types are immune from aliasing.

phase in the full bandwidth may approach or exceed 2.4 rad, where $J_0(m) = 0$. When this happens, all the power goes in the sidebands and the carrier disappears. This is known as the ‘carrier collapse’ in multiplication.

1.5.2 Phase-type (φ -Type) and time-type (x-Type) PM noise

A relevant aspect of the noise behavior of components and systems in the presence of large frequency changes is well described by the terms *phase-type* and *time-type* phase noise (Fig. 1.12). We first introduced this idea in [Calosso and Rubiola, 2017, Sec. 3] to describe noise in the clock distribution of digital electronics. The same terminology applies to

- Amplifiers logic gates, etc., where the output frequency is the same as the input,
- Frequency multiplication, division and synthesis, where the input frequency is scaled up/down by a rational number.
- Frequency synthesis, where the output frequency is set to different values.

Phase-type (φ -type) phase noise

Fig. 1.12a. The quantity $S_\varphi(f)$ is a parameter of the device or system, not affected by the carrier frequency. The dominant noise processes impacts $S_\varphi(f)$ independently of ν_0 , and $S_x(f)$ follows from (6.2). Additive noise in amplifiers (Sec. 16.2), originating from the noise-to-carrier ratio (1.30), is like this. Often flicker PM in amplifier is also this type, described by (1.31). Phase-type noise is also observed in frequency dividers, DDSs and DACs at low ν_o (output), where the scaled-down input noise is lower than the limit set by the output stage.

For example, the white PM noise of a RF amplifier in a reference condition is $S_\varphi(f) = b_0 = 10^{-16} \text{ rad}^2/\text{Hz}$, limited by noise figure and carrier power. Sec. 16.2 explains how this can be calculated. Such noise is of the phase type, shown in Fig. 1.12a. Using (6.2), we get $S_x = 2.53 \times 10^{-32} \text{ s}^2/\text{Hz}$ when the carrier frequency takes the value $\nu_1 = 10 \text{ MHz}$, and $S_x(f) = 1.62 \times 10^{-34} \text{ s}^2/\text{Hz}$ at $\nu_2 = 125 \text{ MHz}$. At ν_2 , the time fluctuation $x(t)$ is a factor $\nu_1/\nu_2 = 0.08$ smaller than at ν_1 , and its PSD is a factor $(\nu_1/\nu_2)^2 = 0.0064$ smaller.

Time-type (x -type) phase noise

Fig. 1.12b. The quantity $S_x(f)$ is a parameter of the device or system, not affected by the carrier frequency. The dominant noise processes impacts on the time fluctuation $S_x(f)$ independently of ν_0 , and $S_\varphi(f)$ follows from (6.2). This behavior is observed in ADCs, DACs and DDSs at high frequencies, where the PM noise is dominated by the jitter of the internal clock distribution. Acoustic/seismic noise and other environment effects on coaxial cables and optical fibers are obviously like this.

For example, the clock distribution of a Cyclone III FPGA has a flicker $S_x = k_{-1} = 4 \times 10^{-28} \text{ s}/\text{Hz}$. Otherway stated, $\sqrt{k_{-1}} = 20 \text{ fs}$. Such noise originates from the time fluctuation in the pipeline of logic gates, which is of the time type (Fig. 1.12b). Using (6.2), we expect a flicker $b_{-1} = 1.6 \times 10^{-10} \text{ rad}^2$ (-98 dBrad^2) at 100 MHz clock, $6.3 \times 10^{-10} \text{ rad}^2$ (-92 dBrad^2) at 200 MHz, and $2.5 \times 10^{-9} \text{ rad}^2$ (-86 dBrad^2) at 400 MHz. The real experiment behind this example [Calosso and Rubiola, 2017, Fig. 6] shows a rather small discrepancy from the $x(\nu_2/\nu_1)^2$ law, $\approx 0.5 \text{ dB}$ for b_{-1} at the three clock frequencies.

A PLL frequency multiplier is another example, where the output is generated by a VCO phase-locked to a 10 MHz reference after frequency division. The division modulo sets the output frequency. The PLL's own noise—not accounting for the external reference—is typically determined by the 10 MHz phase comparator. Assuming that the comparator's flicker is $b_{-1} = 3.2 \times 10^{-10} \text{ rad}^2$ (-95 dBrad^2), the output b_{-1} is of -95 dBrad^2 at 10 MHz output, -89 dBrad^2 at 20 MHz, -83 dBrad^2 at 40 MHz, etc., in 6 dB steps per factor-of-two. This is exactly what we observe on the internal PLL of a Cyclone III FPGA, up to -59 dBrad^2 at 640 MHz, the maximum output frequency [Calosso and Rubiola, 2017, Fig. 6].

Aliasing

The two noise types described may be modified by aliasing, which changes the scaling rule from $(\nu_2/\nu_1)^2$ to ν_2/ν_1 . Aliasing only affects white noise, not flicker and other reddish noise types because the tails of the PSD are too low for the folded-down aliases




Figure 1.13: PN of a Cyclone III, shows φ_T and x_T noise. Courtesy of Claudio E. Calosso

to be significant. Aliasing is typically seen in digital circuits, and in clock distribution of FPGAs and other complex devices as well, at low frequency [Calosso and Rubiola, 2017, Sec. 3].

1.6 Phenomenology of the phase-noise processes

A set of plots showing $\varphi(t)$ and $x(t)$, like Fig. 1.14.

1.7 Statistics and uncertainty

1.8 Amplitude and phase modulation

This should go somewhere else

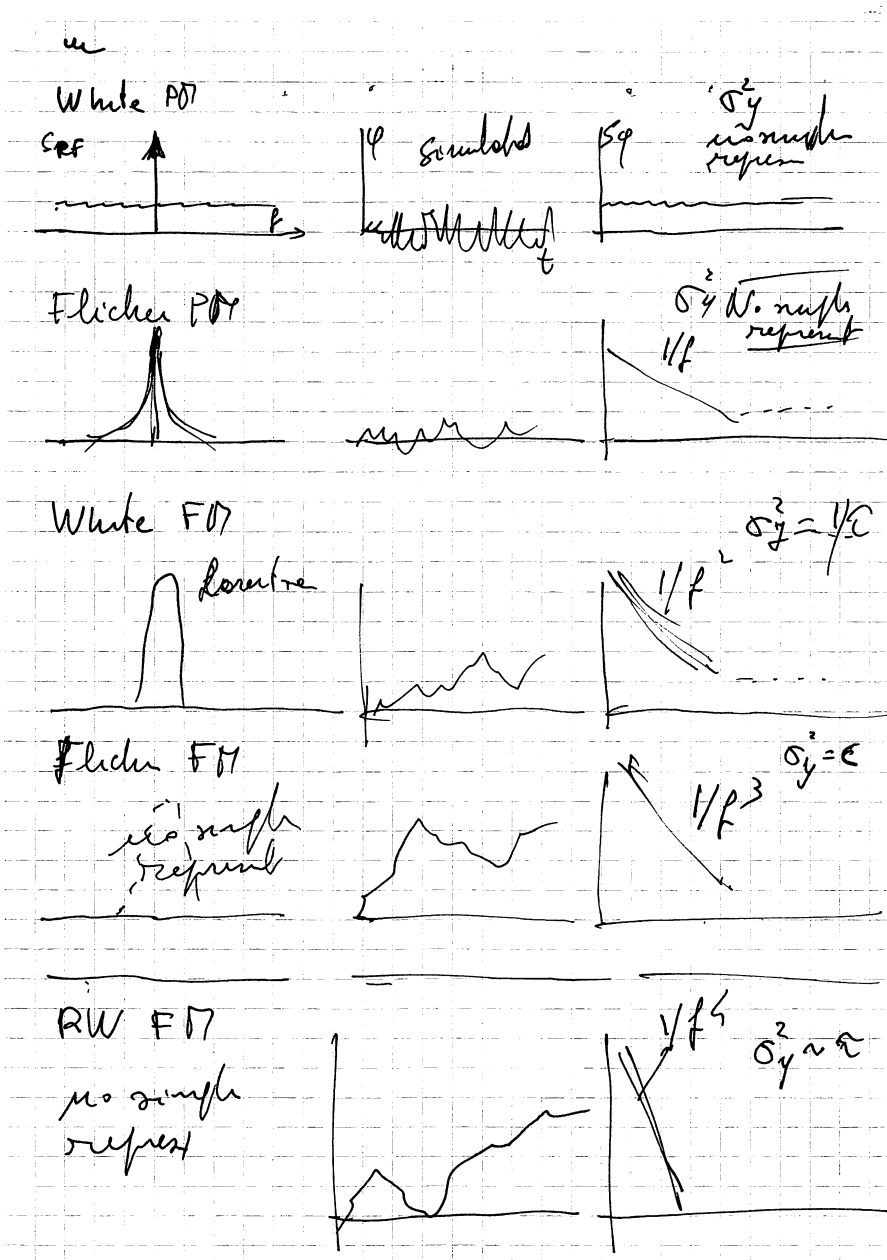


Figure 1.14: .

1.8.1 Sinusoidal modulation

1.8.2 Random modulation (noise)

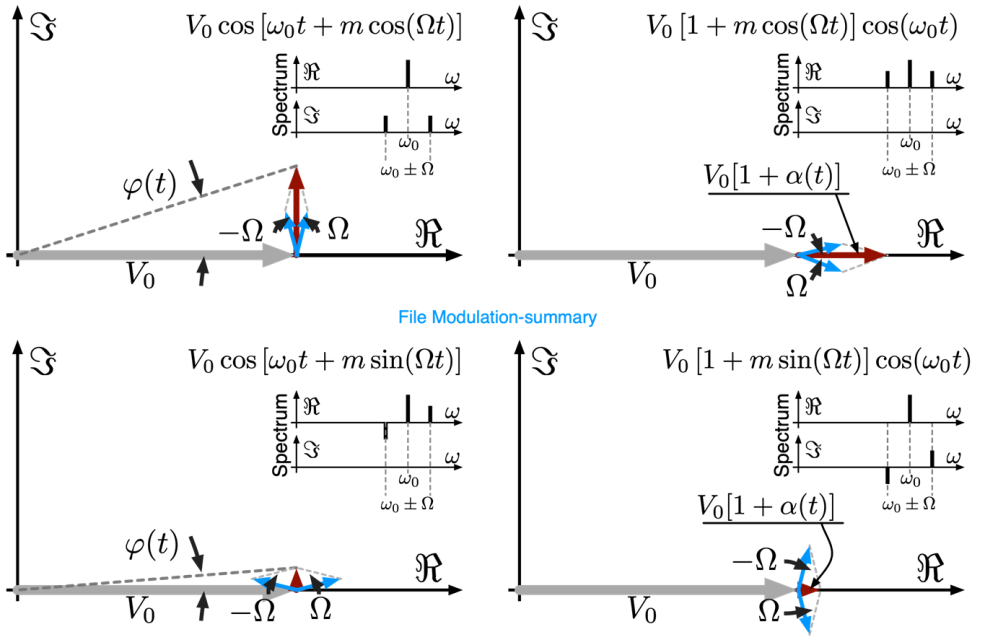


Figure 1.15: Sinusoidal modulations.

1.9 Suggested readings

Digging in the old literature, the earliest use of the phase noise PSD is arguably Victor 1956 [Victor, 1956]. In the NASA Symposium [Chi, 1964] both $S_\varphi(f)$ and the sideband-to-carrier ratio were used, the latter mainly for radar and receivers. However, the symbol $\mathcal{L}(f)$ had still to come. Its early occurrences in archival documents are Blair [Shoaf, 1974] and Shoaf, [Blair, 1974, p. 167 and p. 179–180], both in 1974. Interestingly, the experimental techniques described in [Chi, 1964] rely on phase detectors, which points to $S_\varphi(f)$. Horn [Horn, 1969] is the one and only remarkable counter-example we found, where the sideband-to-carrier ratio is measured with practically useful dynamic range. Note that the history of phase noise is beyond our scope and our skills; we recommend Leeson’s review article [Leeson, 2016, Sec. III–V] and a blog page on Engineering and Technology History Wiki¹³ submitted by Leeson, which refers to additional material of historical interest.

Most of the references about phase noise relate to experimental methods, thus they are moved to Sec. 16.8. The early ideas found in [Chi, 1964] were better formalized in [Barnes et al., 1971] (see also [Rutman, 1978]).

There is a Special Issue of the Proceedings of the IEEE [Chi, 1966], three useful books, Robins 1984 [Robins, 1984], Da Dalt 2018 [Da Dalt and Sheikoleslami, 2018] and Li [Li, 2008], and a free booklet from the NPL [Owen, 2004]. There is very little about AM noise. The reader should refer to [Rubiola, 2005a].

¹³D.B Leeson, *First-Hand: Phase Noise*, https://ethw.org/First-Hand:Phase_Noise.

Chapter 2

Random Noise

Rewrite everything.

Start from B. M. Oliver, Fundamental and quantum noise, 1956

- Thermal noise and noise radiated by a resistor
- Shot noise, electrons and photons. Impulse regime and Gaussian regime, depending on B
- Uncertainty principle and quantum noise in receivers.

Flicker noise

Continue with technical issues

- Noise temperature
- Noise factor and noise figure
- The Rothe Dahlke model

2.1 Thermal (Johnson) noise

Thermal noise is a manifestation of the blackbody radiation which appears as a random fluctuation of the voltage across a device, or a random fluctuation of the current flowing through it. Albeit simpler than the Planck's law of black-body radiation, thermal noise was discovered later. The discovery is generally accredited to John B. Johnson for the experimental observation [Johnson, 1928], and to Harry Nyquist [Nyquist,

1928] for the theoretical explanation. The two articles sit one aside the other in the same issue of the Physical Review in 1928.

The thermal noise can be seen as the energy radiated by a gas of electrons powered by thermal energy. Thermodynamics suggests that the energy is $E = (1/2)kT$ per degree of freedom, thus kT in an electric system which has two degrees of freedom, electric and magnetic field, or equivalently, two orthogonal polarizations states. Because each photon is generated by an accelerated electron, its energy cannot be greater than that of the electron.

The Planck law applied to thermal noise states that the power spectral density, that is, the average available power per unit of bandwidth, delivered by a circuit at temperature T is given by

$$S(\nu) = \frac{h\nu}{\exp\left(\frac{h\nu}{kT}\right) - 1} \quad \text{W/Hz} \quad (\text{Planck law}). \quad (2.1)$$

However, the *receiver* experiences quantum fluctuations $h\nu$, thus the detected signal is seen as

$$S(\nu) = h\nu + \frac{h\nu}{\exp\left(\frac{h\nu}{kT}\right) - 1} \quad \text{W/Hz} \quad (2.2)$$

This is explained in [Oliver, 1965], Sec. ‘Noise in linear amplifier’ on p. 452 (quantum noise), and Sec. ‘Total noise’ on p. 454, where thermal and quantum noise are combined. Equation (2.2) identifies two regions of the spectrum. The *Thermal regime* describes the behavior at low frequencies, where the thermal energy is dominant over the photon energy, $kT \gg h\nu$. The approximation $\exp(h\nu/kT) \simeq 1 + h\nu/kT$ holds, thus the thermal noise is written as.

$$S(\nu) = kT \quad (\text{thermal noise}) \quad (2.3)$$

Conversely, the *Photon regime* describes the behavior at high frequencies, where the photon energy is dominant, $h\nu \gg kT$, and the power spectral density is approximated by

$$S(\nu) = h\nu \quad (\text{photon energy}). \quad (2.4)$$

The two regimes are separated by the cutoff frequency

$$\nu_c = \frac{kT}{h} \ln(2) \quad (\text{cutoff frequency}) \quad (2.5)$$

which results from

$$h\nu = \frac{h\nu}{\exp(h\nu/kT) - 1}$$

Table 2.1 shows the cutoff frequency of (2.2) for some useful temperatures. We notice that true quantum optics experiments are possible above 60 GHz in a liquid-He cryostat, and in the VHF region with a dilution refrigerator. Such experiments are

Table 2.1: Thermal noise and cutoff frequency.

Condition	T	E	Cutoff	Noise, $R = 50 \Omega$
Room temperature	300 K	4.14×10^{-21} J	4.33 THz $69.2 \mu\text{m}$	910 pV/ $\sqrt{\text{Hz}}$ 18.2 pA/ $\sqrt{\text{Hz}}$
T_0 (electronics)	290 K	4.00×10^{-21} J	4.19 THz $71.6 \mu\text{m}$	895 pV/ $\sqrt{\text{Hz}}$ 17.9 pA/ $\sqrt{\text{Hz}}$
Dry ice (-78.5°C)	194.7 K	2.69×10^{-21} J	2.81 THz $107 \mu\text{m}$	733 pV/ $\sqrt{\text{Hz}}$ 17.9 pA/ $\sqrt{\text{Hz}}$
Liquid N ₂	77 K	1.06×10^{-21} J	1.11 THz $270 \mu\text{m}$	461 pV/ $\sqrt{\text{Hz}}$ 9.22 pA/ $\sqrt{\text{Hz}}$
Liquid He	4.2 K	5.80×10^{-23} J	60.7 GHz 4.94 mm	108 pV/ $\sqrt{\text{Hz}}$ 2.15 pA/ $\sqrt{\text{Hz}}$
$^3\text{He}/^4\text{He}$ cryostat	10 mK	1.38×10^{-25} J	144 MHz 2.08 m	5.25 pV/ $\sqrt{\text{Hz}}$ 101 fA/ $\sqrt{\text{Hz}}$

of paramount importance in the emerging domain of quantum computing. A remarkable early example is the demonstration of the Hanbury Brown Twiss effect [Hanbury Brown and Twiss, 1956]¹ at 1.7 GHz by Christian Glattli et al² [Gabelli et al., 2004].

2.1.1 The Thévenin and Norton models for thermal noise

We restrict our attention to the thermal regime. The available power is the maximum power that can be extracted from the circuit, which occurs when the circuit is impedance matched to a load. Letting a resistor R open, no energy is extracted from it, and a voltage fluctuation is present across its ends. Similarly, when the resistor is in short circuit, a current fluctuation flows through it. The power spectral densities are

$$S_v = 4kTR \quad (\text{open-circuit voltage PSD}) \quad (2.6)$$

$$S_i = 4kT/R \quad (\text{short-circuit current PSD}). \quad (2.7)$$

It is a common practice to model a resistor R at temperature T as a cold resistor R in series to a voltage noise generator e_n (Thévenin model) or in parallel to a current noise generator i_n (Norton model), as shown in Fig. 2.1. The statistical nature of voltage and current makes no difference with respect to the dc and sinusoidal regimes freshmen are used to. The quantities e_n and i_n , which represents the rms thermal emf and the rms thermal short-circuit current in 1 Hz bandwidth, are the square root of the respective PSDs

$$e_n \equiv \sqrt{S_v} = \sqrt{4kTR} \quad \text{V}/\sqrt{\text{Hz}} \quad (\text{open circuit voltage}) \quad (2.8)$$

$$i_n \equiv \sqrt{S_i} = \sqrt{4kT/R} \quad \text{A}/\sqrt{\text{Hz}} \quad (\text{short circuit current}) \quad (2.9)$$

¹The original article about the Hanbury Brown Twiss effect.

²Glattli et al, PRL 2004.




Figure 2.1: Thevenin and Norton equivalent circuits for the thermal noise of a resistor.

Albeit the notation e_n and i_n is common in the jargon of electric engineers, the PSD should be preferred.

When the resistor at temperature T is loaded to a cold resistor R_L , power flows from the resistor to the cold load. The maximum power transfer occurs when the load is impedance matched to the generator, $R_L = R$. In this condition, it holds that $V = (1/2)e_n$ and $I = e_n/(2R)$ for the Thévenin model, and $V = (1/2)Ri_n$ and $I = (1/2)i_n$ for the Norton equivalent. The quantity $(1/4)e_n i_n = kT$ is the available power in 1 Hz bandwidth.

2.1.2 Energy exchanged by two resistors at different temperature

The case of two resistors at different temperature³ follows immediately from Fig. 2.1, just attaching a generator to each resistor. The circuit is linear, and the two generators are statistically independent. Introducing the subscripts ‘1’ and ‘2’ for left and right, respectively, we first calculate the power transfer $S_{1,2}$ from left to right with R_1 at temperature T_1 and R_2 at $T = 0$; and $S_{2,1}$ with R_1 at $T = 0$ and R_2 at temperature T_2 . The net result is a PSD $S = k(T_1 - T_2)$ transferred $1 \rightarrow 2$.

Exercise. Two $50\ \Omega$ resistors, R_1 at room temperature (300 K) and R_2 at liquid-N₂ temperature (77 K) are connected through a loss-free coaxial cable having 40 GHz maximum frequency. Everything is impedance matched, and you can assume that the bandwidth limitation is equivalent to a low-pass filter. Calculate the power transferred to the liquid-N₂. (Ans 123 pW).

2.1.3 Thermal noise of an attenuator

We analyze the thermal noise contribution of an attenuator when it is introduced in the signal path. This finds numerous application, among which we mention the loss of antennas and coaxial cables, and the propagation of electromagnetic




Figure 2.2: Thermal noise of an attenuator.

³Need a figure for this

waves through a lossy medium.

Using the model shown in Fig. 2.2, the attenuator has ‘voltage gain’ A , thus power gain $A^2 < 1$ (actually, a loss), and temperature T_a . The attenuator is connected to a source having equivalent temperature T_i , and it is matched to a resistance R at both ends.

The attenuator, obviously, attenuates the noise present at its input as it would do with any other signal. The contribution of the input resistor as seen at the output is

$$A^2 k T_i.$$

Now we temporarily set attenuator and input resistor at the same temperature T . Thanks to impedance matching, the attenuator output is seen as a resistor R at temperature T , thus the noise transferred to the output is kT . We split the latter as

$$kT = A^2 kT + (1 - A^2) kT, \quad \forall A.$$

Because there is no reason for the factors A^2 and $1 - A^2$ to depend on temperature, in the general case the noise seen at the output is

$$kT = A^2 k T_i + (1 - A^2) k T_a,$$

and the contribution of the attenuator is

$$(1 - A^2) k T_a. \quad (2.10)$$

Example (cosmic microwave background). Arno A. Penzias and Robert W. Wilson (Nobel in Physics, 1978) measured the cosmic background cosmic radiation (CMB) with the Horn antenna in Holmdel, NJ, shown aside. They got close to the current estimate, which is a radiation temperature of 2.725 K. Assuming that the temperature of the antenna is $T_a = 290$ K, which is the dissipative loss A of the antenna that gives a noise contribution equal to the temperature of the CMB?




Photo NASA, public domain.

Starting from (2.10), we solve $(1 - A^2) k T_a = k T_{\text{CMB}}$. The result is $A^2 = 1 - T_{\text{CMB}}/T_a = 0.99$ (0.41 dB loss).

2.1.4 Thermal noise across a capacitor

Thermal energy originates a voltage fluctuation across a capacitor (Fig. 2.3) is given by

$$\langle V^2 \rangle = \frac{kT}{C} \quad (2.11)$$

independent of the resistance R . This is relevant in numerous practical cases, among which we mention the gate of CMOS devices, the capacitance of Sample and Hold

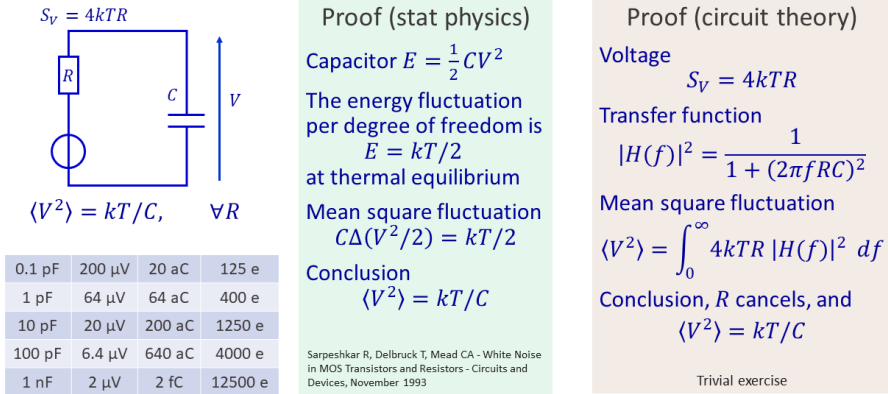


Figure 2.3: Figure Thermal noise across a capacitor.

circuits, and coaxial cables connected to the high-impedance input of instruments like the oscilloscope.

The result is easily proved by feeding the voltage PSD $S_v = 4kTR$ in the filter transfer function

$$|H(f)|^2 = \frac{1}{1 + (2\pi fRC)^2} \quad (2.12)$$

and integrating

$$\langle V^2 \rangle = \int_0^\infty 4kTR |H(f)|^2 df \quad (2.13)$$

It worth mentioning that the resistance R , explicit in the hypothesis, cancel in the calculus.

Another proof, found in reference⁴ starts from the fact that the energy stored in the capacitor is $E = \frac{1}{2}CV^2$, but the voltage across the capacitor is one degree of freedom, thus the energy of the thermal fluctuation is $E = \frac{1}{2}kT$. Combining these two expressions, we need to replace $V^2 \rightarrow \langle V^2 \rangle$, thus $C\langle V^2 \rangle = kT$, and finally $\langle V^2 \rangle = kT/C$.

Table 2.2 gives some reference values for the thermal fluctuation in a capacitor, expressed as a voltage or an electrical charge. Interestingly, the standard capacitors based on the Thompson Lampard theorem used as working standards in primary laboratories, are of the order of 1 pF.

2.2 Shot (Schottky) noise

Shot noise originates from the discrete nature of particles in a transport phenomenon, typically observed in electric circuits and in the detection of photons. Its discovery

⁴Sarpeshkar R, Delbruck T, Mead CA - White Noise in MOS Transistors and Resistors – Circuits and Devices, November 1993

Table 2.2: Thermal noise across a capacitor

0.1 pF	200 μ V	20 aC	125 e
1 pF	64 μ V	64 aC	400 e
10 pF	20 μ V	200 aC	1250 e
100 pF	6.4 μ V	640 aC	4000 e
1 nF	2 μ V	2 fC	12500 e

is generally credited to Schottky⁵. In electrical circuits, shot noise is observed in junctions, vacuum tubes, and other devices or physical experiments where electrons and holes appear as individual particles. In contrast, electrical current in regular conductors, like wires and resistors, is a field, which does not originate shot noise.

The standard picture for the electrical current in vacuum is a stream of ϕ electrons per second emitted at random time, with no memory. The average current is $\langle I \rangle = \phi e$, where $e = 1.602 \times 10^{-19} C$ is the charge of the electron. Unlike thermal noise, shot noise results from a transport phenomenon, thus it is only present when electrical current flows, and it is independent of temperature and of the resistance of the electrical circuit. This type of noise has uniform (white) PSD

$$S_I(f) = 2q \langle I \rangle \quad A^2/\text{Hz} \quad (2.14)$$

The same formula holds in photodiodes, where $I = \eta P/h\nu$ as the photocurrent, ν is the optical frequency, and η is the quantum efficiency, i.e., the probability that a photon is captured and generates an electron-hole pair.

The bandwidth cannot be infinite. Shot noise rolls off at the cutoff frequency

$$f_{\text{sh}} = \frac{1}{2}\phi = \frac{\langle I \rangle}{2q} \quad (2.15)$$

The reason is that the generalized power $S_I(f) f_{\text{sh}}$ must be equal to the generalized dc power $\langle I^2 \rangle$. This is a consequence of the properties of the Poisson distribution. An alternate interpretation relies on the sampling theorem, which states that the maximum frequency of the signal is half the sampling frequency. The same holds for random sampling, just with a smooth cutoff. In this picture, the electrons are electrical pulses playing the role of the random samples.

The shot noise equals the thermal noise at the critical current

$$\langle I \rangle = \frac{kT}{2qR} \quad (2.16)$$

determined by $2q \langle I \rangle R = kT$. For reference, this critical current is of 250 μ A at 290 K with $R = 50 \Omega$, and the associated power is of 3.1 μ W (−25 dBm).

2.2.1 A Deeper sought on shot noise

⁵W. Schottky, Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern, Annalen der Physik 362(23) p541-567, 1918 (in German)

The exponential distribution

A cell emitting particles at random, at the average rate of λ events/s

Probability Density Function

PDF $p(t; \lambda) = \lambda e^{-\lambda t}, t \geq 0$ $\mu = \int t p(t; \lambda) dt = 1/\lambda$

Mean $\mu = 1/\lambda$ $\sigma^2 = \int (t - \mu)^2 p(t; \lambda) dt = 1/\lambda^2$

Variance $\sigma^2 = 1/\lambda^2$

Properties

Memoryless $\mathbb{P}\{T > s + t | T > s\} = \mathbb{P}\{T > t\}$ T is the waiting time

Statistically, T is the same starting at 0 or at s , if the particle did not show up

Maximum differential entropy \rightarrow maximum entropy for a given μ

Featured reading:

W. Feller, *Introduction to probability theory and its applications*, 2nd ed, Wiley. Vol.I, 1957, vol.II, 1970

Vol. 1, Sec. XVII-6 provides a proof that in a memory-less process, the tail of the distribution has to be of the form $u = \exp -\lambda t$ (or zero), and nothing else. See also vol.II, Sec. I-3

Figure 2.4

Shot noise

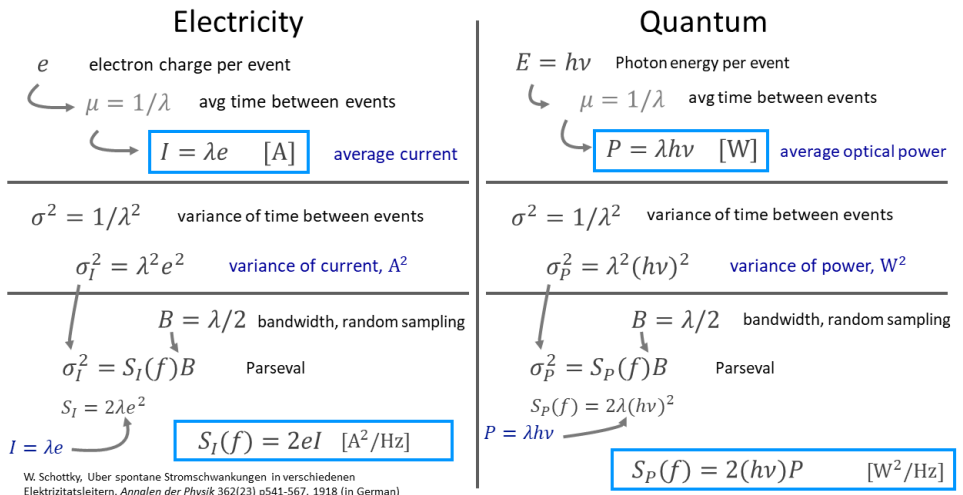


Figure 2.5

2.3 Quantum noise⁶

⁶This Section is inspired to Chapter 2 of E. O. Göbel, U. Siegner, *The New International System of Units*, Wiley VCH 2019 (strongly recommended reading). Please consider also (i) V. B. Braginsky, F. Ya. Khalili, *Quantum Measurement*, Cambridge 1992, and (ii) M. Gläser, M. Kochsiek (Ed.),

Homogeneous Poisson process

An ensemble of memoryless and statistically independent cells emitting at random at the average rate (flux) of ϕ events/s

$$\mathbb{P}\{N(t) = n\} = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$$

\mathbb{P} is the probability that the number N of particles emitted from time 0 to t equals n

This describes “emissions” in physics

- Electrons and holes in a junction
- Photons
- Radioactive decay (assuming that the nuclei are not lost)

Properties

average

$$\mathbb{E}\{N(t)\} = \lambda t \quad \text{written as} \quad \boxed{\mu = \lambda t}$$

variance

$$\mathbb{E}\{[N(t) - \mu]^2\} = \lambda t \quad \boxed{\sigma^2 = \lambda t}$$

signal-to-noise ratio


$$SNR = \sigma / \mu \quad \boxed{SNR = \sqrt{N}}$$

physical meaning of λ

$$\lim_{t \rightarrow \infty} \frac{N(t)}{t} = \lambda \quad \begin{array}{l} \text{average no of events / time,} \\ \text{flux in the case of particle emission} \end{array}$$

W. Feller, Introduction to Probability Theory and Its Applications, vol.II, 2nd ed., Wiley 1970

Figure 2.6



This figure is from
Siebert, B.R.L. and Sommer, K.D.
(2010) in *Uncertainty in Handbook of Metrology*, vol. 2 (eds M. Glaser and M. Kochanski), Wiley-VCH Verlag GmbH, Weinheim, pp. 415–462.
© Wiley-VCH
ISBN 978-3-527-40666-1

Figure 2.7: Thermal and quantum noise

Quantum noise is a very small form of fluctuation, which in our everyday life is deeply hidden under other forms of fluctuations like thermal and shot noise. Nonetheless, it has an important role in our understanding of the physical world.

The Heisenberg principle states that observer, inherently, interacts with the object of the observation extracting an amount of action H which is quantized and multiple of the minimum amount

$$H \gtrsim h .$$

The most usual form of the Heisenberg principle, as commonly found in textbooks, is

$$\Delta p \Delta x \gtrsim \frac{1}{2} \hbar \quad (2.17)$$

where Δp is the uncertainty of the momentum of a particle, and Δx is the uncertainty of position. In electrical measurements, it is convenient to see the action as

$$H = E\tau \text{ or equivalently } H = E \frac{2\pi}{B} \quad (2.18)$$

where E is the energy extracted during the measurement, τ is the duration of the measurement, and $B = 1/2\pi\tau$ is the measurement bandwidth in Hz. Accounting for the minimum action, the energy extracted is

$$E \gtrsim \frac{\hbar}{\tau} \text{ or } E \gtrsim \hbar \frac{B}{2\pi} \quad (2.19)$$

Let us apply this simple concept to the simplest forms of electrical circuit, the resistor, the capacitor and the inductor.

Resistor

The power dissipated by the resistor R is $P = V^2/R = RI^2$. The energy dissipated in the time τ is $E = P\tau$, thus the $E\tau \gtrsim \hbar$ gives

$$\frac{V^2}{R} \tau \gtrsim \frac{\hbar}{\tau} \text{ or } RI^2 \tau \gtrsim \frac{\hbar}{\tau} \quad (2.20)$$

Therefore, the minimum measurable voltage and current are given by

$$V \gtrsim \frac{1}{\tau} \sqrt{\hbar R} \quad (2.21)$$

$$I \gtrsim \frac{1}{\tau} \sqrt{\hbar/R} \quad (2.22)$$

For example, with $\tau = 1$ s, we find 1.8×10^{-16} V and 3.7×10^{-18} A with $R = 50 \, \Omega$; and 2.6×10^{-14} V and 2.6×10^{-20} A with $R = 1 \, \text{M}\Omega$.

Capacitor

The energy stored in the capacitor C is $E = \frac{1}{2} CV^2$, or equivalently $E = \frac{1}{2} \frac{Q^2}{C}$, where Q is the electrical charge. It follows from $E\tau \gtrsim \hbar$ that $\frac{1}{2} CV^2 \gtrsim \frac{\hbar}{\tau}$, and $\frac{1}{2} \frac{Q^2}{C} \gtrsim \frac{\hbar}{\tau}$, thus the minimum measurable voltage and charge are

$$V \gtrsim \sqrt{2\hbar/\tau C}$$

$$Q \gtrsim \sqrt{2\hbar C/\tau}$$

Inductor

The energy stored in the inductor L is $E = (1/2)LI^2$, or equivalently $E = (1/2)(\Phi^2/L)$, where Φ is the magnetic flux. Using $E\tau \gtrsim \hbar$, the minimum measurable




Figure 2.8: A noise spectrum containing flicker. ©Enrico Rubiola, CC-BY-ND 4.0.

current and flux are

$$I \gtrsim \sqrt{2h/\tau L} \quad (2.23)$$

$$\Phi \gtrsim \sqrt{2hL/\tau} \quad (2.24)$$

The equality $E\tau = h$, thus $\frac{1}{2}LI^2 = h/\tau$ occurs with

$$I_0 = \frac{e}{\tau} \quad \text{and} \quad L_0 = \frac{2h}{e^2} \quad (2.25)$$

Similarly, the equality occurs with


$$\Phi_0 = \frac{h}{2e} \quad \text{and} \quad L'_0 = \frac{\tau h}{8e^2} \quad (2.26)$$

Notice that an inductor is quite a good option to measure a magnetic flux, but maybe not the first choice to measure a current.

2.4 Flicker (1/f) noise

Flicker noise is characterized by the PSD $S(f) \propto 1/f^\alpha$, where the exponent α equals one, or is close to one in a wide range of frequency (Fig. 2.8). Actually, α spans from $0.8 \dots 1.2$ to $0.5 \dots 1.5$, depending on the author and on the domain of knowledge. In electronic devices, we observed a nearly exact $1/f$ behavior over 6–8 decades of frequency⁷.

The flicker noise was arguably discovered by Johnson in 1925 observing an anomaly of the Schottky effect in vacuum tubes at low frequencies⁸, and mathematically described by Schottky one year later [Schottky, 1926], before the thermal noise was understood. It was later observed in 1936 by Christiansen as the ‘excess’ noise in carbon microphones [Christensen and Pearson,



⁷Enrico and Claudio Calosso.

⁸J. B. Johnson, The Schottky Effect in Low Frequency Circuits, Phys Rev 26(71), 1 July 1925.




Figure 2.9: Figure A random distribution of low-pass processes in a log scale results in flicker noise.

1936]. The term ‘excess’ means in excess of thermal noise, which at that time was well understood. The most interesting fact about flicker is its ubiquity⁹ [Milotti, 2002, Levitin et al., 2012]. It has been found in geophysical phenomena, climatology, mechanics, optics, classical music, Internet traffic, and in a variety of other domains, and of course electronics. Some cases are well known, but there is still no unifying theory. However, it is clear that a random distribution of low-pass processes on a log-log scale gives rise to a $1/f$ envelope (Figure 12). See Reference [Milotti, 2002] for a review.

Albeit a weak noise phenomenon, flicker is a major issue in different domains like spectral analysis, cryogenic nanodevices and quantum computing¹⁰. A reason is that the resolution cannot be improved by increasing the measurement time. In microwave signals, flicker still originates around DC and modulates the RF. Flicker is of paramount importance for us because it turns out to be a major limitation in the noise of synthesizers, and of oscillators as well.

How small is flicker?

Let us consider a normalized flicker process, $S(f) = 1/f$. The variance of such process results from the trite integral

$$\sigma^2 = \int_a^b S(f) df = \int_a^b \frac{1}{f} df = \ln \left(\frac{b}{a} \right)$$

where a and b are the low-pass and high-pass cutoff frequencies. Albeit σ^2 diverges for $a \rightarrow 0$ or $b \rightarrow \infty$, it remains surprisingly small for any realistic value of a and b . To convince the reader, we take $a = 2.3 \times 10^{-18}$ Hz and $b = 2.95 \times 10^{42}$ Hz as, arguably, the largest possible integration interval. The lower bound is the reciprocal of the age of the universe, $A_U = 4.35 \times 10^{17}$ s (13.8 billion years), and the upper bound is $1/2\pi\tau_P$, where $\tau_P = \sqrt{\hbar G/c^5} \simeq 5.39 \times 10^{-44}$ s is the Planck time. The result is

$$\sigma^2 = \ln \left(\frac{1/2\pi\tau_P}{1/A_U} \right) \simeq 138.4 \quad (21.4 \text{ dB}) \quad (2.27)$$

So, if an operational amplifier having a flicker of $4.5 \text{ nV}/\sqrt{\text{Hz}}$ at $f = 1 \text{ Hz}$ (the popular low-noise OP27/37) was created at the time of the big bang and, and if it

⁹Milotti 2002, Levitin, Chordia, & Menon, 2012.

¹⁰References!!!




Figure 2.10: Integrated $1/f^\alpha$ noise. The black (lower) curve refers to the bandwidth of an FFT analyzer (1 mHz to 10 MHz). The blue (middle) curve refers to an experiment limited by the professional life of a researcher (30 years), up to the frequency of optical light. The orange (upper) curve refers to the cosmological/fundamental limits stated in the text. ©Enrico Rubiola, CC-BY-ND 4.0.

could deliver noise up to $b = 2.95 \times 10^{42}$ Hz with full impunity, the flicker noise integrated over such bandwidth would be $\sigma = 53 \text{ nV}/\sqrt{\text{Hz}}$.

That said, $\sigma^2 = \int (1/f^\alpha) df$ grows larger at low frequencies if $\alpha > 1$, or at higher frequencies if $\alpha < 1$. Figure 2.10 shows the total noise power for different values of α in three asymptotic scenarios: (i) a table-top experiment, (ii) a lifetime project, and (iii) cosmological/fundamental limits.

Flicker in mechanical systems


Flicker in ferromagnetic systems

Hysteresis is identified as a mechanism which turns thermal noise (white) into $1/f$ noise. In ferromagnetic materials, this is known as the Barkhausen effect, and the noise generated in this way is sometimes called Barkhausen noise. In mechanical systems, the combined effect of thermal energy and hysteresis is a good candidate to limit the dimensional stability of a solid. Understanding this mechanism has applications in metrology, where the mechanical stability of Fabry-Pérot cavities limits the frequency stability of lasers.

Let us start with the Debye-Einstein theory of heat capacity in solids. The solid is made of atoms kept in place by electrical forces, assimilated to damped mass-spring oscillators¹¹. Such oscillators are coupled, thus for a large number of atoms there is a continuum of permitted vibrating frequencies (Fig. 2.11). The energy equipartition

¹¹Damped oscillators, called resonators in the jargon of engineering.

Flicker ($1/f$) dimensional fluctuation is powered by thermal energy



A single theory explains

- Heat capacity
- Thermal expansion
- Elasticity
- ... and their fluctuations




Figure 2.11: fig:Debye-Einstein

suggests that atoms vibrate powered by thermal energy, uniformly distributed over frequency. Damped oscillators release energy to the thermal bath, getting colder than the average temperature, and in turn they re-absorb energy from the thermal bath. This continuous exchange is the soul of the fluctuation dissipation theorem. Looking at the nature of dissipation, we realize that it cannot be viscous. In fact, the viscous dissipation is characterized by force proportional to velocity. The consequence is that the matter, rather being solid, would flow down slowly under the effect of gravity¹², and ultimately flatten on the lowest horizontal surface available.


$$\ddot{v}_o + \frac{\omega_n}{Q} \dot{v}_o + \omega_n^2 v_o = \frac{\omega_n}{Q} \dot{v}_i \quad (2.28)$$

Flicker in semiconductors (Hossein Taheri, 2023)

Regarding the following / attached slide on flicker noise, I came across a justification in a YouTube video for $1/f$ noise in the context of electronic semiconductors:

Burst or popcorn noise, related to the capture and release of charge carriers in semiconductor defects, has an autocorrelation function of the form $R - xx(\tau) = R_{xx}(0) \exp(-|\tau|/\tau)$, such that its Fourier transform results in a power spectral density of the form $S_x(\omega) = 4R_{xx}(0)\tau/[1 + (\omega\tau)^2]$, which corresponds to the $1/f^2$ curves in your slides. Here τ is a constant indicating the trapping time of the carrier in the defect. A combination of many such events where the probability of the trapping $N(\tau) = B/\tau$ (B a constant) decreases with increasing trapping time τ (i.e., the extreme events of very long trapping time are very rare, while very short trapping

¹²The reader may have heard that the glasses of some ancient cathedrals tend to become thinner on the top and thicker on the bottom, suggesting that they flow down like a very viscous liquid. However, this effect is way too slow for being the result of viscosity.



There is no viscous dissipation in solids

Dissipation is structural (hysteresis)

Structural dissipation

micro/nanoscale, instantaneous

Dissipated energy $E = \int F dx$

Small vibrations

The hysteresis cycle keeps the aspect ratio

$$E \propto x_0^2 \quad \text{Energy lost in a cycle}$$

Thermal equilibrium

$$P = kT \quad \text{in 1 Hz BW}$$

$$P \propto kT x_0^2$$

$$x_0^2 \propto 1/f \rightarrow \text{flicker}$$

Figure 2.12: fig:Structural-dissipation

times are quite common) can be found by the following integral

$$S_y(\omega) = \int_0^\infty N(\tau) \frac{4R_{xx}(0)\tau}{1 + \omega^2\tau^2} d\omega = \int_0^\infty \frac{B}{\tau} \frac{4R_{xx}(0)\tau}{1 + \omega^2\tau^2} d\omega \quad (2.29)$$

This integral is of the form $\int du/(1 + u^2)$, with $u = \omega\tau$, which is the arctangent function. So, the overall result is that $S_y(\omega) \propto 1/\omega$, i.e., similar to $1/f$ noise.

Bibliography of flicker noise

C. J. Christiansen, G. L. Pearson, Spontaneous Resistance Fluctuations in Carbon Microphones and Other Granular Resistances, BSTJ 15(2) p.197-223, April 1937. Arguably, the discovery of flicker.

F. N. Hooge, $1/f$ noise is no surface effect, Phys Lett 29(3) p.139-140, 21 April 1969. Classical article.

D. J. Levitin, P. Chordia, V. Menon, Musical Rythm Spectra from Bach to Joplin Obey to $1/f$ Power Law, Proc. Nat. Academy of Science 109(10) p.716-3720, February 2012.

A. L. Mcwhorter, $1/f$ Noise and Germanium Surface Properties, Proc. Semiconductor Surface Physics p.2017-228, June 1956. Classical article.

E. Milotti, $1/f$ noise, a pedagogical review, arXiv.physics 0204033, April 2002.

Paladino E et al., $1/f$ Noise, Implications for solid-state quantum information, Rev Modern Phys 86, April-June 2014.

Numata K, Kemery A, Camp J, Thermal-noise limit in the frequency stabilization of lasers with rigid cavities, Phys Rev Lett 93(25) 250602, December 2004.

P. R. Saulson, Thermal Noise in Mechanical Experiments, Phys Rev D 42(8), October 1990.

A. van der Ziel, Unified Presentation of $1/f$ Noise in Electronic Devices: Fundamental $1/f$ sources, Proc IEEE 76(3), March 1988.

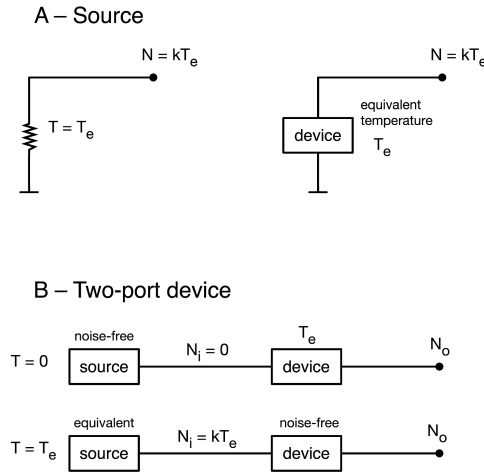


Figure 2.13: Noise temperature (Wiley, needs permission)

L. K. J. Vandamme, G. A. Trefan, A review of $1/f$ noise in bipolar transistors, *Fluct Noise Lett* 1(4) 2001.

M. B. Weissman, $1/f$ noise and other slow, nonexponential kinetics in condensed matter, *Rev Modern Phys* 60(2) p.537-571, April 1988.

2.5 Technical representations of noise

2.5.1 The equivalent noise temperature

The noise equivalent temperature, denoted with T_e , is another widely used parameter to assess the noise, by analogy with the thermal noise.

The case of a source (Fig. fig:Noise-temperature-A) is quite straightforward. We say that the source has equivalent noise temperature T_e at a specific frequency if the available noise power at its output in 1 Hz bandwidth is $N = kT_e$, the same of a resistor at temperature T_e .

The case of a two-port device is similar, and illustrated in Fig. 2.13-B. The available output noise N_o is the same for the two cases, (1) when the real device is connected to a noise-free source, and (2) when a noise-free, and otherwise equal device is connected to a source at the temperature T_e .

2.5.2 Noise factor and noise figure

Noise Factor, Noise Figure, and Noise Temperature

The noise factor, or equivalently the noise figure, is likely the most used, if not over-used, parameter to characterize the noise of two-port components or systems, like

amplifiers, frequency converters, radio receivers etc. However often used interchangeably, one should prefer the term noise factor for the dimensionless quantity F , and noise figure (NF) for F expressed in dB

$$\text{NF} = 10 \log_{10} F \quad [\text{dB}] \quad (2.30)$$

A popular definition of F , given by Friis [Friis, 1944] in 1944, is

$$F = \frac{\text{SNR}_i}{\text{SNR}_o} \quad (\text{Friis, obsolete definition}) \quad (2.31)$$

where SNR is the Signal-to-Noise Ratio, and the subscripts “ i ” and “ o ” stand for input and output. A substantially equivalent, and less known definition was given by North¹³ in 1942 as the ratio of (1) the output noise power from the transducer to (2) the output noise power from an equivalent noise-free transducer.

The problem with Eq. (2.30), and also with the North’s definition, is that the degradation to the SNR depends on the noise of the source that excites the device. This is seen by sending a signal of power P_i and thermal noise kT_i to a device of power gain A^2 . Denoting with N_{dev} the available noise contribution of the device in 1 Hz bandwidth, as observed at the output, Eq. (2.30) gives

$$F = \frac{P_i}{kT_i} \frac{N_{\text{dev}} + A^2 kT_i}{A^2 P_i} = 1 + \frac{N_{\text{dev}}/A^2}{kT_i} \quad (2.32)$$

For the same device at the same temperature, N_{dev} is the same, but F is affected by the temperature T of the input termination.

The ambiguity of Eq. (2.30) is solved with the new definition proposed by the IRE¹⁴, and adopted by the NIST¹⁵ (at that time, NBS)

Definition (IRE): The noise factor, at a specified input frequency, is defined as the ratio of (1) the total noise power per unit bandwidth at a corresponding output frequency available at the output port when the noise temperature of the input termination is standard (290 K) to (2) that portion of (1) engendered at the input frequency by the input

It is worth mentioning that the IRE definition differs from the Friis definition only in the use of the reference temperature $T_0 = 290\text{K}$. For this reason, Eq. (2.30) is still found in the literature, with the reference temperature often implied. Both definitions are smart in that the input circuit is taken in its actual configuration, making no assumption on impedance matching and on best noise impedance, if different from the characteristic impedance.

The noise temperature is related to the noise factor. Assuming impedance matching at the input, the device’s noise contribution observed at the output is $A^2 kT_e$. Accordingly, the noise factor (IRE definition) can be written as

$$F = \frac{A^2 kT_e + A^2 kT_0}{A^2 kT_e} \quad (2.33)$$

¹³North, 1942

¹⁴Haus 6011033, look in the original Word file.

¹⁵NBS adopted the new definition, look in the original Word file.

Noise power vs R

40




Figure 2.14: Noise power vs R . Move this figure to an appropriate place.

thus

$$F = 1 + \frac{T_e}{T_0} \quad (2.34)$$

As we mentioned, the SNR degradation produced by a two-port device depends on the input noise, hence on the temperature of the source. Figure 2.15 shows an example. All the plots match the Friis' definition $F = \text{SNR}_i / \text{SNR}_o$, but only the plot $T_i = 290\text{K}$ matches the IRE definition.

The noise factor of an attenuator is an interesting case because of the physical insight it provides. We first assume that the attenuator is impedance-matched at both ends, and that it receives the thermal noise kT_i from a resistor. The attenuator 'amplifies' the input noise by a factor of $A^2 < 1$, as it does with any signal. Yet, the attenuator adds its own contribution, which we will calculate. For this purpose, we set the temperature of both input resistor and attenuator to T_0 . The output noise is the sum of (1) the input noise attenuated, that is, $N' = kT_0 A^2$, and (2) the noise $N'' = A^2 kT_e$ from the attenuator. Because the attenuator output is equivalent to a resistor at the temperature T_0 , the total available noise at the output is equal to $N_0 = kT_0$. Thus, $N'' = kT_0 - N'$, and consequently $N'' = kT_0(1 - A^2)$. Referring the attenuator noise to the input, we find $N_e = N'' / A^2 = kT_0(1 - A^2)/A^2$. Finally, the equivalent temperature is $T_e = T_0(1 - A^2)/A^2$. Using $F = 1 + T_e/T_0$, we find $F = 1/A^2$. This proves the thumb rule that the noise factor of an attenuator is equal to the power attenuation.

The reference temperature $T_0 = 290\text{ K}$ ($17.2\text{ }^\circ\text{C}$) is often used in electronics and radio engineering because the corresponding thermal energy

$$kT_0 = 4 \times 10^{-21} \text{ W/Hz} \quad \text{i.e., } -174 \text{ dBm/Hz} \quad (2.35)$$

is a convenient round number, at a temperature which is decent approximation of the physical temperature found in most practical cases. At T_0 , the equivalent noise




Figure 2.15: SNR degradation due to a two-port device of noise temperature T_e for various values of temperature T_i of the input source. From Fundamentals of RF and Microwave Noise Figure Measurements, September 2017, © 2010-2017 Keysight Technologies (Find the citation in the original Word file). Needs permission, and adapted to the text.

voltage kT_0/q is equal to 25 mV. In the jargon of semiconductors, this quantity is often denoted with V_T . The thermal emf in 1 Hz bandwidth is denoted with $e_n = \sqrt{4kT/R}$. Across a $R = 50 \Omega$ resistor we find

$$e_n = 0.9 \text{ nV}/\sqrt{\text{Hz}}, \quad \text{or} \quad -181 \text{ dB}\sqrt{\text{Hz}} \quad (\text{open circuit}). \quad (2.36)$$

2.5.3 The Rothe Dahlke model

The noise factor and the noise temperature are simplified representations of the reality. More accurate noise models, generally good for virtually all practical purposes, resort to the seminal paper by Rothe and Dahlke¹⁶ on the theory of linear four-poles. Following this approach, the noise is best described by adding a voltage generator e_n and a current generator i_n at the device input, as shown on Figure 2.16. Such generators are described in terms of their power spectral densities $S_e(f)$, $S_i(f)$ and $S_{ie}(f)$, or equivalently in terms of their variances and the covariance $\langle i_n e_n^2 \rangle$ in 1 Hz bandwidth. The most relevant fact is that e_n and i_n define an impedance which generally differs from the input impedance. Thus, the device has two “optimum” impedances, one for maximum power transfer, which refers to the usual conjugate matching, and one for lowest noise. Consequently, the impedance of the generator impacts on the noise factor, and low noise design requires a tradeoff between gain and noise.

¹⁶Rothe Dahlke, the original article.




Figure 2.16: Generalized noisy device (Wiley, needs permission).

After doing the appropriate math, the noise factor is given by

$$F(\Gamma_g) = F_o + 4 \frac{R_n}{Z_0} \frac{|\Gamma_o - \Gamma_g|^2}{|1 + \Gamma_o|^2 (1 - |\Gamma_g|^2)} \quad (2.37)$$

where F_o is the minimum (optimum) noise factor, R_n is the noise resistance (the sensitivity of noise factor to source resistance changes), Z_0 is the nominal input, Γ_g is the reflection coefficient, and Γ_o is the value of Γ_g with which F_o is achieved. The parameters F_o , R_n , and Γ_o describe the noise of the device, and Γ_g is a free design choice. Represented on the Smith chart, the noise factor $F(\Gamma_g)$ looks like a set of equal-noise circles (see Figure 2.17).

Packaged amplifiers are often impedance-matched in a wide range of frequency, thus the noise factor is degraded because impedance matching is often privileged versus noise matching, and also because of the loss of the input circuit. Values of 1 dB to 4 dB are rather common. Conversely, the noise factor of a transistor can be quite low, yet at the cost of uncomfortable impedance matching (see the example on Table 3).

When several stages are cascaded, the noise factor of the chain is given by the Friis formula¹⁷

$$F = F_1 + \frac{F_2 - 1}{A_1^2} + \frac{F_3 - 1}{A_1^2 A_2^2} + \dots \quad (2.38)$$

Accordingly, the first stage of a chain should have a low noise factor, while the noise factor requirement of subsequent stages is relaxed. The ideas underneath the Friis formula are simple. The overall noise factor F is referred to the input of the chain. The noise contribution of each amplifier is added as square voltage, power, or PSD, as appropriate, because the amplifier are separate devices, and their fluctuations are statistically independent. The contribution F_1 of the first stage is obvious. The second stage has no resistive load at the input, thus no kT_0 . There remains $(F_2 - 1)kT_0$, which becomes $(F_2 - 1)kT_0/A_1^2$ referred at the input of the chain, after dividing by the power gain A_1^2 of the first stage. Recursively, the noise $(F_3 - 1)kT_0$ of the third stage is divided by the power gain $A_1^2 A_2^2$ of the two preceding stages, etc.

¹⁷Cite the original article of the Friis formula




Figure 2.17: Example of equal-noise circles of an amplifier, on the Smith chart. Edited, from Fundamentals of RF and Microwave Noise Figure Measurements, September 2017, © 2010-2017 Keysight Technologies (Find the citation in the original Word file). Needs permission and adapted to the text.

Table 2.3: Typical noise parameters of the TAV-581+ transistor.

ν_0 , GHz	F_0	$ \Gamma_0 $, deg	$\arg(\Gamma_0)$, deg	R_n , Ω	Gain, dB
0.5	0.09	0.37	16.1	4.0	26.6
0.7	0.12	0.37	28.5	3.5	24.6
0.9	0.16	0.37	40.6	3.0	23.0
1.0	0.18	0.37	46.6	3.0	22.3
1.9	0.34	0.39	97.4	1.5	17.8
2.0	0.35	0.39	102.7	1.5	17.4
2.4	0.42	0.40	123.4	1.5	16.3
3.0	0.53	0.41	152.5	1.5	14.9
3.9	0.69	0.43	−168.1	2.5	13.3
5.0	0.89	0.45	−127.1	5.0	11.8
5.8	1.03	0.46	−102.1	8.0	10.8
6.0	1.06	0.47	−96.5	9.0	10.6

Conditions: $V_{DS} = 4$ V, $I_{DS} = 30$ mA.
Data are taken from the TAV-581+ data sheet
(Mini Circuits, NY, USA).

The Friis formula is an approximation based on impedance matching. A more accurate model should account for two main facts. First, impedance mismatch calls for a correction term which lowers the gains A_i^2 , based on the reflection coefficients between the i -th and the $(i + 1)$ -th stage. Second, the noise factor F_1 should account




Figure 2.18: Noise in bipolar transistors.

for the source impedance, and likewise all the F_{i+1} should account for the output impedance of the preceding, i -th, stage.

2.6 Noise in selected devices

2.6.1 Bipolar transistors

2.6.2 Field-effect and MOS transistors

2.6.3 The Rothe Dahlke model Noise in four poles


2.6.4 Noise in operational amplifiers

2.7 The Measurement of noise temperature

The equivalent temperature of a device is usually measured with the Y method shown on Figure 2.21. The method consists of switching two impedance-matched input sources, one at the temperature T_h (hot) and the other at the temperature T_l (cold). Asymptotically, if one can set $T_h \rightarrow \infty$ and $T_l \rightarrow 0$, the temperature T_h is the probe signal that enables the measurement of the power gain A^2 , and T_l gives the equivalent temperature T_e after taking away the gain A^2 . In actual experimental conditions, the output noise power is

$$P_h = A^2 k(T_h + T_e)B$$

$$P_l = A^2 k(T_c + T_e)B$$




Noise is modeled as a voltage generator $e(t)$ and a current generator $i(t)$

Consequences

- The golden rule $Z_L = Z_G^*$ is broken
- Three different impedance-matching criteria at Port 1 (the device is the load)
 - Lowest noise: $Z_G = e_n/i_n$
 - Maximum power: $Z_L = Z_G$
 - Highest Signal-To-Noise Ratio (SNR): something in between

H. Rothe, W. Dahlke, Theory of Noisy Fourpoles, Proc IRE 44(6) p.811-818, June 1956
H. A. Haus, R. B. Adler, Circuit theory of linear noisy networks, John Wiley & Sons 1959

Figure 2.19: The Rothe Dahlke model.



Need to design precision electronics?

- D. Feucht, Analog Circuit Design Series, 4 volumes, SciTech 2010, ISBN 978-1-891121-XY-Z (old school but great)
- S. Franco S, Design with operational amplifiers and analog integrated circuits 4ed, McGraw Hill 2015, ISBN 978-0-07-802816-8 (best for designing with operational amplifiers)
- P. Horowitz, W. Hill, The Art of Electronics 3ed, Cambridge 2015, ISBN 978-0-521-80926-9 (Bible of instrument design, physical insight)
- Tietze U, Schenk C, Gamme E - Electronic Circuits 2ed - Springer 2007, ISBN 978-3-540-78655-9

$$a \oplus b = (1/a + 1/b)^{-1}$$

Noise resistance

$$R_{eq} = R_p + (R_N \oplus R_F)$$

Voltage

$$V = V_{OS} + R_p I_p - (R_N \oplus R_F) I_N$$

Split I_N and I_p into offset and bias, $I_{OS} \pm \frac{1}{2} I_B$

$$\text{Bias } I_B = \frac{1}{2} (I_p - I_N), \text{ Offset } I_{OS} = I_p - I_N$$

Total effect

$$V = V_{OS} + \frac{1}{2} [R_p - (R_F \oplus R_N)] I_B + [R_p + (R_N \oplus R_F)] I_{OS}$$

Obvious extension to noise

$$V^2 = \sum_i V_i^2$$

Figure 2.20: Noise in operational amplifiers.




Figure 2.21: The Y method for the measurement of the equivalent noise temperature of a device (Wiley, needs permission).

where B is the bandwidth of a filter at the device output. The solution of the system is

$$T_e = \frac{T_h - Y T_l}{Y - 1} \quad \text{with} \quad Y = \frac{P_h}{P_l} = \frac{T_h + T_e}{T_l + T_e} > 1 \quad (2.39)$$

The main virtue of the Y method is that the factor A^2/B cancels in the evaluation of Y . This simplifies the calibration and results in improved accuracy because A generally suffers from flatness defect, while B is the equivalent noise bandwidth, which results from integrating the transfer function.

It is worth mentioning that the equivalent noise temperature includes thermal noise in strict sense, the shot noise, and any other noise process. For this reason, people with a background in optics may find this concept particularly misleading. In fact, in optical systems there is no temperature and, in high SNR condition, the electrical noise at the detector output is chiefly shot noise.

2.8 Points of interest

2.8.1 Blackbody radiation

Every physical body above the absolute zero temperature emits energy, which is described by the spectral radiance

$$B(\nu, T) = \frac{2h\nu^3}{c^2} \frac{1}{\exp\left(\frac{h\nu}{kT}\right) - 1} \quad [\text{Wsr}^{-1}\text{m}^{-2}\text{Hz}^{-1}]. \quad (2.40)$$

The spectral radiance is the power emitted at the frequency ν by a body at the temperature T , normalized to a surface of 1 m^2 , a solid angle of 1 sr , and 1 Hz bandwidth. The blackbody radiation was already measured accurately at the end of 1800, but it was Max Planck in 1900 the first who found a satisfactory explanation in 1900 using empirical constants and assuming the quantization of energy. Later, he proved that this is necessary for thermodynamic equilibrium.

A different perspective is given by Larmor, who calculated the energy radiated by an accelerated particle, directed in the plane orthogonal to the acceleration vector. Blackbody radiation results from the random acceleration of the “gas” of electrical charges in thermodynamic equilibrium. Indeed, the Larmor’s work adds the coherent emission. Antennas radiate coherently because of the acceleration inherent in oscillating current.

Strictly speaking, electrical charges are even not necessary for a body to emit energy because accelerated magnetic dipoles radiate too. The blackbody radiation of neutron stars and other exotic cosmic objects can be interpreted in this way.


2.8.2 The dilution refrigerator

The dilution refrigerator enables extremely low temperatures, down to 2–100 mK depending on volume and isolation of the cryostat. The theory was formulated by Heinz London in the early 1950s, and it was implemented in 1964 at the Kamerlingh

- ^4He is a boson
 - Superfluidity
- ^3He is a fermion
 - Pauli exclusion principle
 - Fermi liquid
- Pump (mainly) ^3He
- Push the fluid into the unstable region
- Phase separation is endothermal

Theory: Heinz London, early 1950s
 Implementation: 1964, Kamerlingh Onnes Lab, Leiden
 H. K. Onnes (Nobel 1913) liquefied He (1908) and discovered the superconductivity of Hg (1911)

Featured reading:
 Chapter 9, S. W. Van Sciver, Helium cryogenics 2nd ed., Springer 2012



The dilution refrigerator at the FEMTO-ST Institute

Figure 2.22: Dilution refrigerator: phase diagram (left), and a photo of the Oxford Instruments refrigerator at FEMTO-ST Institute (right).

Onnes Lab in Leiden, The Netherlands. It is now a commercial product, available from Oxford Instruments and Bluefors, among others.

The dilution refrigerator exploits the strange properties of a mixture of liquid ^3He (rare isotope) and ^4He (common Helium) at low temperature. The former is a fermion (half-integers spin number), thus is governed by the Pauli exclusion principle. The latter is a boson (integer spin). Below 2.2 K, superfluidity shows up in ^4He , so the mixture is a ^3He Fermi liquid (non-interacting particles) in superfluid ^4He . Below 870 mK, the mixture is unstable, and tends to split into two separate phases. Such splitting is the endothermal reaction exploit to refrigerate the mixture.

2.8.3 A short review of the Harry Nyquist's article

This Section needs partial rewriting.

Harry Nyquist gave a remarkably smart explanation to the Johnson's observation, at once heuristic, rigorous and concise. We review the Nyquist's digression, which starts from the Johnson's experimental observation of current fluctuations in resistors.

Connecting two equal resistors at the same temperature as in Fig. 2.23-1, the current fluctuations of one resistor results in a dissipated power in the other resistor, and vice versa. However, thermodynamics suggests that the power flowing in one direction is equal to the power flowing in the other direction, regardless of the nature of such resistors. If a loss-free filter is inserted between the resistors — Nyquist used a notch (Fig. 2.23-2), but a modern engineer would rather prefer a passband — the thermodynamic equilibrium still holds. The preliminary conclusion is that the thermal emf across a resistor is a function of frequency, resistance and temperature only.

The next step targets at identifying the power spectral density of the thermal emf. This is accomplished by introducing an ideal electrical line between the two resistors (Fig. 2.23-3). Such ideal line does not loose energy, for example in the form of dissipation or radiation, and is impedance matched, thus its characteristic



Image user Qulbik, Wikimedia

The Harry Nyquist's Article

Thermal equilibrium




Fig. 1.

Thermal equilibrium also applies to any frequency (interval)
EMF E is a function of R , T and f only




Fig. 2.

Loss-free, terminated electrical line




Fig. 3.

After thermal equilibrium, isolate the line (short at both ends).

Modes at $\nu = n c / \ell$

ν = frequency, c = velocity

Energy kT per mode

$$dE = 2\ell kT d\nu / c$$

Average power in frequency $d\nu$,
and in time ℓ/c is $kT d\nu$

Extension to electrical circuits




Fig. 4.

Energy per degree of freedom

$$h\nu / (e^{h\nu/kT} - 1)$$

instead of kT

Conclusion

$$E_V^2 d\nu = 4R_V h d\nu / (e^{h\nu/kT} - 1)$$

J. B. Johnson, Thermal Agitation of Electricity in Conductors, Phys Rev 32(1) p.97-109, July 1928
H. Nyquist, Thermal agitation of electric charges in conductors, Phys Rev 32(1) p.110-113, July 1928

Figure 2.23: The Nyquist's interpretation of thermal noise.

impedance is equal to R . Let us denote the line length with l , and the speed of propagation with v . When the thermal equilibrium is established, the same amount of power flows in either direction, which is represented as waves in the line. Then, the line is shorted at both ends, so that the energy trapped in the line is completely reflected and bounces infinitely. The line is a resonator oscillating at its natural frequencies $v/2l$, $2v/2l$, $3v/2l$, etc. For large l , the number of modes in the frequency interval $d\nu$ is $2l d\nu/v$. The average energy per mode is kT , half in the electric field and half in the magnetic field. Thus, the total energy in $d\nu$ is $2l kT d\nu/v$. This is the energy that was transferred from the resistors to the line during the time of transit l/v . Consequently, the average power transferred from each resistor to the line in the frequency interval $d\nu$ during the time interval l/v is $kT d\nu$. This is the same as the PSD $S = kT$. Accounting for the Thévenin model, the PSD of the thermal EMF is $S_v = 4kTR$.

2.8.4 Noise and entropy

2.8.5 Thermal noise in optical fibers

Jing Dong et al., Observation of fundamental thermal noise in optical fibers down to infrasonic frequencies Appl. Phys. Lett. 108, 021108 (2016). DOI 10.1063/1.4939918

Chapter 3

Power Spectra

Physical concept of PSD

The PSD is the distribution of power vs. frequency (power in 1-Hz bandwidth)

The PS is the distribution of energy vs. frequency (energy in 1-Hz bandwidth)

Power (energy) in physics is a square (integrated) quantity

PSD W/Hz, or V²/Hz, A²/Hz, rad²/Hz etc.

$$S_v(f) = \frac{\langle v_B^2(f) \rangle}{B} \quad (3.1)$$

The analogy shown in Fig... highlights the important properties of the PSD

Cross spectrum

The noise of a device under test (DUT) is measured simultaneously with two instruments, each of which contributes its own background. The average cross power spectral density converges to the DUT power spectral density. This method enables the extraction of the DUT noise spectrum, even if it is significantly lower than the background. After a snapshot on practical experiments, we go through the statistical theory and the choice of the estimator. A few experimental techniques are described, with reference to phase noise and amplitude noise in RF/microwave systems and in




Figure 3.1: Physical concept of power spectral density.




Figure 3.2: Basics of the cross-spectrum method.

photonic systems. The set of applications of this method is wide. The final section gives a short panorama on radio-astronomy, radiometry, quantum optics, thermometry (fundamental and applied), semiconductor technology, metallurgy, etc.

This report is intended as a tutorial, as opposed to a report on advanced research, yet addressed to a broad readership: technicians, practitioners, Ph.D. students, academics, and full-time scientists.

3.1 Introduction

Measuring a device under test (DUT), the observed spectrum contains the DUT noise, which we can call *signal* because it is the object of the measurement, and the background noise of the instrument. The core of the cross-spectrum measurement method is that we can measure the DUT simultaneously with two equal instruments. Provided that experimental skill and a pinch of good luck guarantee that DUT and instruments are statistically independent, statistics enables to extract the DUT spectrum from the background.

The two-channel measurement can be modeled as the block diagram of Fig. 3.2, where $a(t)$ and $b(t)$ are the background of the two instruments, and $c(t)$ the DUT noise, under the hypothesis that $a(t)$, $b(t)$ and $c(t)$ are statistically independent. Thus, the




Figure 3.3: Average and deviation of the cross spectrum $|\langle S_{yx} \rangle_m|$, as a function of the number m of averaged realizations of white Gaussian noise. Since the statistical properties of $S_{yx}(f)$ are the same at any frequency, only one point (i.e., one frequency) is shown and the variable f is dropped. The DUT noise is 10 dB lower than the background.

observed signals are

$$\begin{aligned} x(t) &= c(t) + a(t) \\ y(t) &= c(t) + b(t) . \end{aligned}$$

We are interested in the power spectral density¹ (PSD), which is a normalized form of spectrum that expresses the power per unit of bandwidth, denoted with $S(f)$. It will be shown that the average cross-PSD $\langle S_{yx}(f) \rangle$ converges to the DUT PSD $S_{cc}(f)$, which is what we want to measure.

The idea of the cross-spectrum method is explained in Fig. 3.3. This figure builds from the output of the free-running analyzer, after selecting one frequency (f_0). This is a sequence of $|S_{yx}(f_0)|$ called realizations, which we average on contiguous groups of m values $|\langle S_{yx}(f) \rangle_m|$. The averages form a (slower) sequence whose statistical properties depend on m . So, Fig. 3.3 plots the average and the variance of the sequence of averages, as a function of m . At small values of m , the background is dominant and decreases as m increases. Beyond $m \approx 100$, we observe that $|\langle S_{yx}(f) \rangle_m|$ stops decreasing and approaches the value of 0.1 (−10 dB), which is the DUT noise in this example. The standard deviation further decreases. The background is dominant below $m \approx 100$. Beyond, the DUT noise shows up and the estimation accuracy increases, as seen from the deviation-to-average ratio. Notice that the choice of $|\langle S_{yx}(f) \rangle_m|$ as an estimator of $S_{yx}(f)$ is still arbitrary and will be further discussed.

All this report is about how and why the cross-spectrum converges to the DUT noise $S_{cc}(f)$, and about how this fact can be used in the laboratory practice. The scheme of Fig. 3.2 is analyzed from the following standpoints

¹The PSD as a statistical concept will be defined afterwards. Newcomers can provisionally use $S_{yx}(f) = \frac{1}{T} Y(f) X^*(f)$, which is the readout of the FFT analyzer. T is the measurement time.

Normal use. All the noise processes $[a(t), b(t) \text{ and } c(t)]$ have non-negligible power. We use the statistics to extract $S_{cc}(f)$.

Statistical limit. In the absence of correlated phenomenon, thus with $c = 0$, the average cross spectrum takes a finite nonzero value, limited by the number of averaged realizations.

Hardware limit. After removing the DUT, a (small) correlated part remain. This phenomenon, due to crosstalk or to other effects, limits the instrument sensitivity.

Though the author is inclined to use phase and amplitude noise as the favorite examples (Section 3.8.1 and 3.8.2), the cross-spectrum method is of far more general interest. Examples from a variety of research fields will be discussed in Section 3.8.3.

As a complement to this report, the reader is encouraged to refer to classical textbooks of probability and statistics, among which [?, ?, ?] are preferred.

3.2 Power spectral density

The processes we describe are stationary and ergodic. The requirement that noise be stationary and ergodic is not a stringent constraint in the laboratory practice because the words ‘stationary’ and ‘ergodic’ are the equivalent of ‘repeatable’ and ‘reproducible’ in experimental physics. Thus, a realization $x(t)$ has the same statistical properties independently of the origin of time, and also the statistical properties of the entire process $\mathbf{x}(t)$. Unless otherwise specified, $\mathbf{x}(t)$ is a zero-mean finite-power process. The power spectral density (PSD) of such processes is

$$S_{xx}(f) = \mathcal{F} \{R_{xx}(t')\} \quad (3.2)$$

where $\mathcal{F}\{ \}$ is the Fourier transform operator,

$$R_{xx}(t') = \mathbb{E} \{ \mathbf{x}(t) \mathbf{x}(t + t') \} \quad (3.3)$$

the autocorrelation function, and $\mathbb{E}\{ \}$ the mathematical expectation.

As a simplified notation, we use the upper case for the Fourier transform, and the left-right arrow for the transform inverse-transform pair, thus

$$x(t) \leftrightarrow X(f) \quad \text{Fourier transform – inverse transform pair}.$$

The two-sided Fourier transform and spectra are generally preferred in theoretical issues, while the experimentalist often prefers the single-sided representation. Though we use the one-sided representation in all figures, often we do not need the distinction between one-sided and two-sided representation. In most practical measurements the Fast Fourier Transform (FFT) replaces the traditional Fourier transform, and the frequency is a discrete variable.

The Wiener-Khintchine theorem for ergodic and stationary processes enables to calculate the PSD through the absolute value of the Fourier transform. Thus it holds

that

$$\mathbb{E}\{S_{xx}(f)\} = \mathbb{E}\left\{\lim_{T \rightarrow \infty} \left[\frac{1}{T} X_T(f) X_T^*(f)\right]\right\} \quad (3.4)$$

$$= \mathbb{E}\left\{\lim_{T \rightarrow \infty} \left[\frac{1}{T} |X_T(f)|^2\right]\right\}, \quad (3.5)$$

where the subscript T means truncated over the measurement time T , and the superscript ‘*’ stands for complex conjugate. By the way, the factor $\frac{1}{T}$ is necessary for $S_{xx}(f)$ to have the physical dimension of a *power density*, i.e., power per unit of frequency.

Omitting the expectation, (3.4) can be seen as a realization of the PSD. In actual experiments the expectation is replaced with the average on a suitable number m of spectrum samples

$$\langle S_{xx}(f) \rangle_m = \frac{1}{T} \langle |X_T(f)|^2 \rangle_m \quad (\text{avg, } m \text{ spectra}). \quad (3.6)$$

As an obvious extension, the cross PSD of two generic random processes $\mathbf{x}(t)$ and $\mathbf{y}(t)$

$$S_{yx}(f) = \mathcal{F}\{R_{yx}(t')\} \quad (3.7)$$

is measured as

$$\langle S_{yx}(f) \rangle_m = \frac{1}{T} \langle Y_T(f) X_T^*(f) \rangle_m. \quad (3.8)$$

3.2.1 Measurement time T

In practical experiments the measurement time is finite, so we can only access the truncated version $x_T(t) \leftrightarrow X_T(f)$ of a realization. In order to simplify the notation, the subscript T for the *truncation time* will be omitted. Thus for example we write (3.8) as

$$\langle S_{yx}(f) \rangle_m = \frac{1}{T} \langle Y(f) X^*(f) \rangle_m \quad (\text{abridged notation}).$$

3.2.2 Why white Gaussian noise

However too simplistic at first sight it may seem, the use of white Gaussian noise is justified as follows. First, spectrally-smooth noise phenomena originate from large-number statistics (electrons and holes, semiconductor defects, shot noise, etc.), which by virtue of the central limit theorem yield to Gaussian process. Second, most non-white noise phenomena of interest in follow the power-law model $S(f) = \sum h_i f^i$, hence they can be converted into white noise after multiplication by a suitable power of f without affecting the PDF, and converted back after analysis. The idea of whitening and un-whitening a noise spectrum is by the way of far broader usefulness than shown here. For these reasons we can take full benefit from the simplicity of white Gaussian noise. Yet, it is understood that white noise rolls off at some point, so that all signals have finite power.

3.3 The cross-spectrum method

Recalling the definitions of Section 3.1, we denote with $a(t)$ and $b(t)$ the background of the two instruments, with $c(t)$ the common noise, and with A , B and C their Fourier transform, letting the frequency implied. Working with realizations, we no longer need a separate notation for the process. By definition, $a(t)$, $b(t)$ and $c(t)$ are statistically independent. We also assume that they are ergodic and stationary. The two instrument outputs are

$$x(t) = c(t) + a(t) \quad \leftrightarrow \quad X = C + A \quad (3.9)$$

$$y(t) = c(t) + b(t) \quad \leftrightarrow \quad Y = C + B. \quad (3.10)$$

First, we observe that the cross-spectrum S_{yx} converges to S_{cc} . In fact,

$$\begin{aligned} \mathbb{E}\{S_{yx}\} &= \frac{1}{T} \mathbb{E}\{YX^*\} \\ &= \frac{1}{T} \mathbb{E}\{[C + A] \times [C + B]^*\} \\ &= \frac{1}{T} [\mathbb{E}\{CC^*\} + \mathbb{E}\{CB^*\} + \mathbb{E}\{AC^*\} + \mathbb{E}\{AB^*\}] \\ &= S_{cc} \end{aligned} \quad (3.11)$$

because the hypothesis of statistical independence gives

$$\mathbb{E}\{CB^*\} = 0, \quad \mathbb{E}\{AC^*\} = 0, \quad \text{and} \quad \mathbb{E}\{AB^*\} = 0.$$

Then we replace the expectation with the average on m measured spectra

$$\begin{aligned} \langle S_{yx} \rangle_m &= \frac{1}{T} \langle YX^* \rangle_m \\ &= \frac{1}{T} \langle [C + A] \times [C + B]^* \rangle_m \\ &= \frac{1}{T} [\langle CC^* \rangle_m + \langle CB^* \rangle_m + \langle AC^* \rangle_m + \langle AB^* \rangle_m] \\ &= S_{cc} + O(\sqrt{1/m}), \end{aligned} \quad (3.12)$$

where $O(\)$ means ‘order of.’ Owing to statistical independence, the cross terms decrease proportionally to $1/\sqrt{m}$.

3.3.1 Statistical limit

With no DUT noise it holds that $c = 0$, hence $S_{cc} = 0$. Maintaining the hypothesis of statistical independence of the two channels, we notice that the number of averaged spectra sets a statistical limit to the measurement. Only the cross terms remain in (3.12), which decrease proportionally to $1/\sqrt{m}$. Thus, the statistical limit is

$$\langle S_{yx} \rangle_m = \frac{1}{T} \langle AB^* \rangle_m \approx \sqrt{\frac{1}{m} \langle S_{yy} \rangle_m \langle S_{xx} \rangle_m} \quad (\text{statistical limit}). \quad (3.13)$$

Accordingly, a 5 dB improvement on the single-channel noise costs a factor of 10 in averaging, thus in measurement time. The convergence law will be extensively discussed in Section 3.6.

3.3.2 Hardware limit

Breaking the hypothesis of the statistical independence of the two channels, we are interested in the *correlated noise* of the instrument, which limits the sensitivity. This can be due for example to the crosstalk between the two channels, or to environmental fluctuations (ac magnetic fields, temperature, etc.) acting simultaneously on the two channels. The mathematical description is simplified by setting the true DUT noise to zero, and by re-interpreting $c(t)$ as the *correlated noise* of the instrument observed on unlimited number of averaged spectra

$$\mathbb{E}\{S_{yx}\} = \mathbb{E}\{S_{cc}\} \quad (\text{hardware limit}) . \quad (3.14)$$

Nonetheless, the correct identification of this limit may require non-trivial experimental skill.

3.3.3 Regular DUT measurement

The accurate measurement of a regular DUT requires that

1. The number m is large enough for the statistical limit to be negligible
2. The hardware background noise is negligible as compared to the DUT noise

In this conditions, the average cross spectrum converges to the expectation of the DUT noise

$$\langle S_{yx} \rangle_m \rightarrow \mathbb{E}\{S_{cc}\} \quad (\text{DUT measurement}). \quad (3.15)$$

This is the regular use of the instrument.


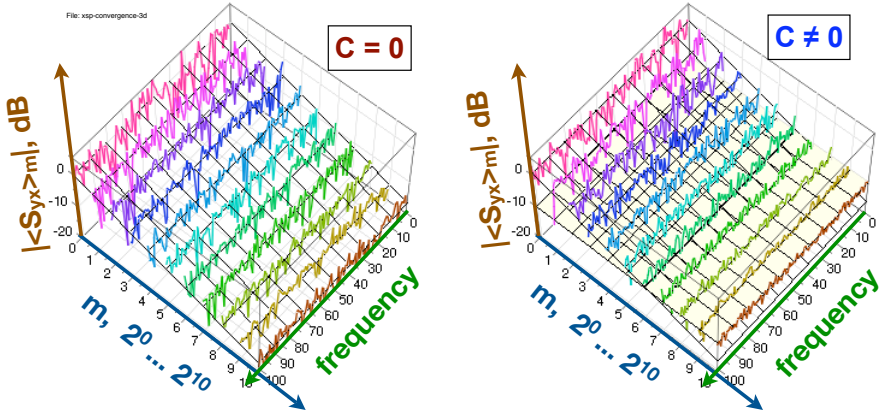
3.4 Running the experiment

Before getting through mathematical details, it is instructive to start from a simplified picture of what happens when we run an experiment. For this purpose, we chose $\hat{S}_{yx} = |\langle S_{yx} \rangle_m|$ as an estimator of S_{yx} , which is often the default of the FFT analyzer in cross-spectrum mode. This estimator is suitable to be displayed on a logarithmic scale (dB) because it takes only nonnegative values, but it is biased. We observe the PSD on the display of the FFT analyzer as m increases, looking for the signature of \hat{S}_{yx} converging to S_{cc} .

We restrict our attention to the case of DUT noise smaller than the single-channel background, as it usually occurs when we need the correlation. The purpose for this assumption is to make the simulations representative of the laboratory practice. And of course we assume that the two channels are equal.

3.4.1 Ergodicity

Averaging on m realizations, the progression of a measurement gives a sequence of spectra $|\langle S_{yx} \rangle_m|_i$ of running index i , as shown in Fig. 3.4. For a given frequency f_0 ,

Figure 3.4: Sequence of cross spectra $|\langle S_{yx}(f) \rangle_{32}|$.Figure 3.5: Sequence of cross spectra $|\langle S_{yx} \rangle_m|$.

the sequence $|\langle S_{yx}(f_0) \rangle_m|_i$ is a time series. Since $S_{yx}(f_1)$ and $S_{yx}(f_2)$, are statistically independent for $f_1 \neq f_2$, also $|\langle S_{yx}(f_1) \rangle_m|_i$ and $|\langle S_{yx}(f_2) \rangle_m|_i$ are statistically independent. For this reason, scanning the frequency axis gives access to (a subset of) the statistical ensemble.

Ergodicity allows to interchange time statistics and ensemble statistics, thus the running index i of the sequence and the frequency f . The important consequence is that the average and the deviation calculated on the frequency axis give access to the average and deviation of the time series, without waiting for multiple realizations to be available. This property helps detect when the cross spectrum leaves the $1/\sqrt{m}$ law and converges to the DUT noise.

Figure 3.5 shows a sequence of cross spectra $|\langle S_{yx} \rangle_m|$, increasing m in powers of two. On the left-hand side of Fig. 3.5, the DUT noise is set to zero. Increasing m ,

the average cross spectrum decreases proportionally to $1/\sqrt{m}$, as emphasized by the slanted plane. The $1/\sqrt{m}$ law is easily seen after averaging on the frequency axis separately for each value of m , and then transposing the law to each point of the frequency axis thanks to ergodicity. The right-hand side of Fig. 3.5 shows the same simulation, yet with the DUT noise set to a value of 10 dB lower than the single-channel background. At small values of m the cross-spectrum is substantially equal to the previous case. Yet at $m \gtrsim 100$ the cross-spectrum leaves the $1/\sqrt{m}$ law (slanted plane) and converges to the DUT noise (horizontal plane at -10 dB). Once again, thanks to ergodicity we can transpose the average on the frequency axis to each point of the frequency axis.

In the rest of this Section we will refer to a generic point of the PSD, letting the frequency unspecified. The variable f is omitted in order to simplify the notation. Hence for example we will write $\Re\{S_{yx}\}$ instead of $\Re\{S_{yx}(f)\}$.

3.4.2 Single-channel noise.

It is explained in Sec. 3.5 that the single-channel PSD $\langle S_{xx} \rangle_m$ is χ^2 distributed with $2m$ degrees of freedom. The average PSD is equal to $\frac{1}{T} \mathbb{V}\{X\} = \frac{1}{T} \mathbb{V}\{A\} + \frac{1}{T} \mathbb{V}\{C\}$, where $\mathbb{V}\{\}$ is the variance. The deviation-to-average ratio is equal to $1/\sqrt{m}$. Of course the same holds for S_{yy} , after replacing A with B .

The track seen on the display converges to the DUT noise *plus* the background noise, and shrinks as m increases. The track thickness is twice the deviation. This fact is shown on Fig. 3.6. The green plot, labeled $|S_{xx}|$, keeps the same vertical position as m increases, and shrinks.

3.4.3 Cross-spectrum observed with insufficient m .

When the number m of averaged realizations is insufficient for the DUT noise to show up, the system behaves as the two channels were (almost) statistically independent. In this conditions we can predict the spectrum by setting $X \simeq A$, $Y \simeq B$ and $C \simeq 0$, thus $\mathbb{E}\{S_{yx}\} \simeq 0$.

Increasing m , the statistical distribution of the estimator $\hat{S}_{yx} = |\langle S_{yx} \rangle_m|$ converges to a Rayleigh distribution with $2m$ degrees of freedom. Normalizing on the single-channel background $\mathbb{E}\{S_{xx}\} = \mathbb{E}\{S_{yy}\} = 1$, and using the results of Sec. 3.6, we find that

$$\begin{aligned} \mathbb{E}\{\hat{S}_{yx}\} &= \mathbb{E}\{|\langle S_{yx} \rangle_m|\} = \sqrt{\frac{\pi}{4m}} = \frac{0.886}{\sqrt{m}} \\ \mathbb{V}\{\hat{S}_{yx}\} &= \mathbb{V}\{|\langle S_{yx} \rangle_m|\} = \frac{1}{m} \left(1 - \frac{\pi}{4}\right) = \frac{0.215}{m}, \end{aligned}$$

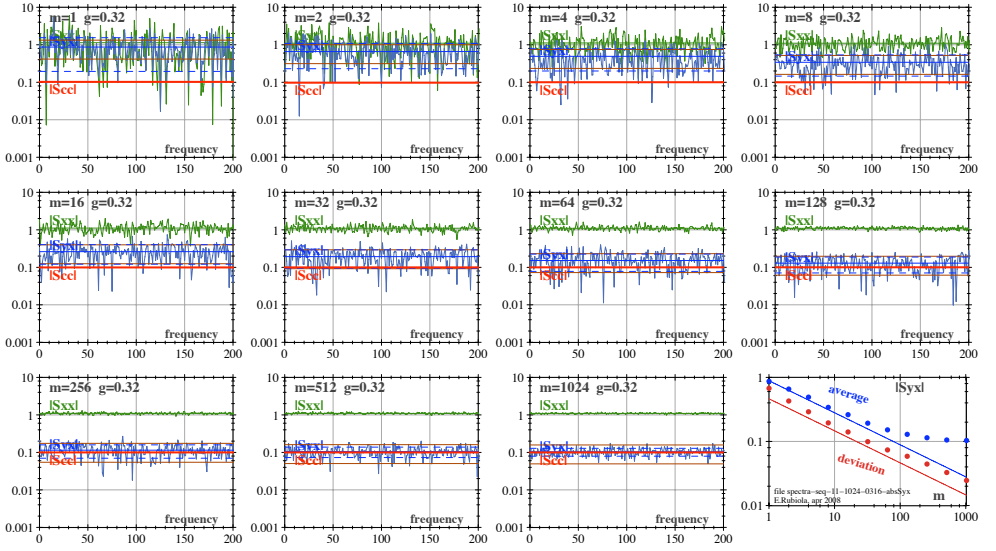


Figure 3.6: Simulated PSD, plotted for increasing number m of averaged realizations. The parameter $g = 0.32$ (-10 dB), which is κ in the main text, is the correlated noise, while the single-channel background is of one.

and therefore

$$\text{dev}\{\hat{S}_{yx}\} = \sqrt{\mathbb{V}\{|\langle S_{yx} \rangle_m|\}} = \sqrt{\frac{1}{m} \left(1 - \frac{\pi}{4}\right)} = \frac{0.463}{\sqrt{m}}$$

$$\frac{\text{dev}\{\hat{S}_{yx}\}}{\mathbb{E}\{\hat{S}_{yx}\}} = \sqrt{\frac{4}{\pi} - 1} = 0.523 \quad (\text{independent of } m).$$

The track is centered at $\frac{0.886}{\sqrt{m}}$. This is the estimator bias. The track looks as a horizontal band located at $\text{avg} \pm \text{dev}$, thus on a logarithmic from $10 \log_{10}(1 - \text{dev}/\text{avg}) = -3.21$ dB to $10 \log_{10}(1 + \text{dev}/\text{avg}) = +1.83$ dB asymmetrically distributed around the average. This is shown on Fig. 3.6. For $m \lesssim 100$, the blue plot labeled $|S_{yx}|$ decreases proportionally to $1/\sqrt{m}$ and has the constant thickness of half a decade (5 dB), independent of m .

3.4.4 Cross-spectrum observed with large m .

When the number m of averaged realizations is large enough, the background noise vanishes and the DUT spectrum shows up. The cross spectrum no longer decreases but the variance still does. Qualitatively speaking, the average is set by the DUT noise S_{cc} and the deviation is set by the instrument background divided by \sqrt{m} . On a logarithmic scale, the track no longer decreases and starts shrinking. This is shown on Fig. 3.6 for $m \gtrsim 100$, blue plot labeled $|S_{yx}|$.

The above reasoning can be reversed. The simultaneous observation that the cross spectrum *stops decreasing*, and *shrinks* is the signature that the averaging process is converging. The single-channel background is rejected and the instrument measures the DUT noise (or the hardware limit, which is higher). This fact is of paramount importance in some measurements, where for some reasons we cannot remove the DUT.

3.5 Estimation of S_{xx}

The measurement accuracy depends on three main factors, instrument calibration, instrument background (front-end and quantization), and statistical estimation. Only the latter is analyzed in this Section.

As a property of zero-mean white Gaussian noise, the Fourier transform $X = X' + \imath X''$ is also zero-mean Gaussian, and the energy is equally split between X' and X'' . Restricting our attention to a generic point (i.e., to an unspecified frequency), the PSD is

$$\mathbb{E}\{S_{xx}\} = \frac{1}{T} \mathbb{E}\{|X|^2\} = \frac{1}{T} \mathbb{E}\{[X'^2 + X''^2]\} .$$

For use in this Section we define

$$\zeta^2 = \mathbb{E}\{S_{xx}\} ,$$

which is the power in 1 Hz bandwidth. Since X' and X'' are zero-mean Gaussian-distributed random variables, a single realization

$$S_{xx} = \frac{1}{T} [X'^2 + X''^2]$$

follows a χ^2 distribution with two degrees of freedom. After our definition of ζ^2 , we find that

$$\mathbb{V}\{X'\} = \mathbb{V}\{X''\} = \frac{T}{2} \zeta^2 .$$

because S_{xx} includes a factor $\frac{1}{T}$. This is seen on the “scaled χ^2 ” column of Table 3.1, after setting $\nu = 2$ (degrees of freedom) and $\sigma = \frac{1}{2}T\zeta^2$. On that Table we find that $\mathbb{E}\{S_{xx}\} = \frac{1}{T}\nu\sigma^2$, which is equal to ζ^2 , and that $\mathbb{V}\{S_{xx}\} = \frac{1}{T^2}2\nu\sigma^4$, hence

$$\mathbb{V}\{S_{xx}\} = \zeta^4 .$$

Averaging on m realizations of S_{xx}

$$\langle S_{xx} \rangle_m = \frac{1}{m} \sum_{i=1}^m \frac{1}{T} [X_i'^2 + X_i''^2] ,$$




Figure 3.7: Probability density function $f(x)$ of the PSD averaged on m realizations.

we notice that $\langle S_{xx} \rangle_m$ has χ^2 distribution with $2m$ degrees of freedom. Using the right-hand column of Table 3.1, we find $\mathbb{V}\{\langle S_{xx} \rangle_m\} = \frac{1}{m}\zeta^4$. The uncertainty (standard deviation) is therefore

$$\text{dev}\{\langle S_{xx} \rangle_m\} = \frac{1}{\sqrt{m}}\zeta^2 \qquad \frac{\text{dev}\{\langle S_{xx} \rangle_m\}}{\mathbb{E}\{\langle S_{xx} \rangle_m\}} = \frac{1}{\sqrt{m}} .$$

Figure 3.7 shows an example PDF of the spectrum averaged on m realizations. The χ^2 distribution is normalized for the standard deviation to be equal one. Increasing m , the PDF converges to the normal distribution and shrinks.

Finally, we may find useful the following normalization

$$S_{aa} = 1 \quad (\text{background}) \qquad S_{cc} = \kappa^2 \quad (\text{DUT}) .$$

Expanding $X = X' + \imath X'' = (A' + C') + \imath(A'' + C'')$ we notice that X is zero-mean white Gaussian noise, and that

$$\mathbb{E}\{\langle S_{xx} \rangle_m\} = 1 + \kappa^2 \qquad \text{dev}\{\langle S_{xx} \rangle_m\} = \frac{1 + \kappa^2}{\sqrt{m}} .$$

3.6 Estimation of S_{yx} and noise rejection

It is obvious from Eq. (3.6) that the spectrum $S_{xx}(f)$ takes always *real positive* values, regardless of the number of averaged realizations. Since some kind of fundamental noise is always present in a physical experiment, the probability that $S_{xx}(f)$ nulls at some frequency is zero. Conversely, the cross-spectrum $S_{yx}(f)$ is a *complex* function that converges to the positive function $S_{cc}(f)$ only after averaging on a sufficient number m of realizations, as seen in Eq. (3.12).

In numerous practical cases we need to plot $S_{yx}(f)$ on a logarithmic vertical scale, which is of course impossible where $S_{yx}(f)$ is not positive.

- In radio engineering virtually all spectra are given in decibels, which resorts to a logarithmic scale.
- When the spectrum spreads over a large dynamic range, only a compressed scale makes sense. The logarithmic scale is by far the preferred representation.
- Numerous spectra found in physical experiments follow a polynomial law because the time-domain derivative (integral) maps into a multiplication (division) of the spectrum by f^2 . On a logarithmic plot, a power of f maps into a straight line.
- It is explained in Section 3.4 that running the experiment, average and deviation of the instrument noise are ruled by the same $1/\sqrt{m}$ law until the number of averaged realizations is sufficient for $S_{yx}(f)$ to converge to $S_{cc}(f)$. This is most comfortably seen on a logarithmic scale.

Thus, we need to extend Section 3.5 to the cross spectrum, discussing the suitable estimators. The estimator may introduce noise and bias. In everyday life a better estimator may save only a little amount of time, and in this case it could be appreciated mainly because it is smarter. Oppositely in long-term measurements, like *timekeeping* and *radioastronomy*, a single data point takes years of observation. Here, the choice of the estimator may determine whether the experiment is feasible or not.

3.6.1 Basic material

Let us expand S_{yx}

$$\begin{aligned}
 S_{yx} &= \frac{1}{T} \mathbb{E} \{YX^*\} \\
 &= \frac{1}{T} \mathbb{E} \{(B+C) \times (A+C)^*\} \\
 &= \frac{1}{T} \mathbb{E} \{(B' + \imath B'' + C' + \imath C'') \times (A' - \imath A'' + C' - \imath C'')\} \\
 &= \frac{1}{T} \mathbb{E} \{(B'A' + B''A'' + B'C' + B''C'' + C'A' + C''A'' + C'^2 + C''^2) \\
 &\quad + \imath(B''A' - B'A'' + B''C' - B'C'' + C''A' - C'A'')\} \tag{3.16}
 \end{aligned}$$

and simplify the calculus by normalizing on the variances as follows

$$\begin{array}{lll}
 \mathbb{V}\{A\} = 1 & \mathbb{V}\{A'\} = 1/2 & \mathbb{V}\{A''\} = 1/2 \\
 \mathbb{V}\{B\} = 1 & \mathbb{V}\{B'\} = 1/2 & \mathbb{V}\{B''\} = 1/2 \\
 \mathbb{V}\{C\} = \kappa^2 \ll 1 & \mathbb{V}\{C'\} = \kappa^2/2 & \mathbb{V}\{C''\} = \kappa^2/2 .
 \end{array}$$

Notice that an additional factor T must be added a-posteriori for a proper normalization on $\mathbb{E}\{S_{aa}\} = \mathbb{E}\{S_{bb}\} = 1$ (background power in 1 Hz bandwidth equal to one), as we did in Section 3.5. Thanks to energy equipartition, it follows that $\mathbb{V}\{A'\} = 1/2 \Rightarrow \mathbb{V}\{A'\} = T/2$, etc.

The assumption that $\kappa^2 \ll 1$, though not necessary, is quite representative of actual experiments because the main virtue of the correlation method is the capability of extracting the DUT noise when it is lower than the background.

Looking at (3.16), we identify the following classes

terms	\mathbb{E}	\mathbb{V}	PDF	comment
$B'A', B''A'', B''A', B'A''$	0	1/4	$\frac{1}{2\pi} K_0\left(\frac{ x }{2}\right)$	product of zero-mean Gaussian processes
$B'C', B''C'', C'A', C''A'', B''C', B'C'', C''A', C'A''$	0	$\kappa^2/4$	$\frac{1}{2\pi\kappa} K_0\left(\frac{ x }{2\kappa}\right)$	product of zero-mean Gaussian processes
$C'^2 + C''^2$	κ^2	κ^4	$\chi^2_{\nu=2}$	sum of zero-mean square Gaussian processes

Equation (3.16) can be rewritten as

$$S_{yx} = \frac{1}{T} \mathbb{E} \{ \mathcal{A} + \imath \mathcal{B} + \mathcal{C} \} \quad (3.17)$$

where the terms

$$\begin{aligned} \mathcal{A} &= B'A' + B''A'' + B'C' + B''C'' + C'A' + C''A'' \\ \mathcal{B} &= B''A' - B'A'' + B''C' - B'C'' + C''A' - C'A'' \\ \mathcal{C} &= C'^2 + C''^2 \end{aligned}$$

have the statistical properties listed underneath. By virtue of the central limit theorem, the probability density function of $\langle \mathcal{A} \rangle_m$ and of $\langle \mathcal{B} \rangle_m$ converge rapidly to a Gaussian for large m . Notice that $\langle \mathcal{C} \rangle_m$ follows a χ^2 distribution with $2m$ degrees of freedom, thus for large m it can be approximated with a Gaussian distributed variable of equal average and variance, which is denoted with $\langle \tilde{\mathcal{C}} \rangle_m$.

term	\mathbb{E}	\mathbb{V}	PDF	comment
$\langle \mathcal{A} \rangle_m$	0	$\frac{1 + 2\kappa^2}{2m}$	Gauss	large- m average (sum) of zero-mean Gaussian processes
$\langle \mathcal{B} \rangle_m$	0	$\frac{1 + 2\kappa^2}{2m}$	Gauss	
$\langle \mathcal{C} \rangle_m$	κ^2	κ^4/m	$\chi^2_{\nu=2m}$	average (sum) of chi-square processes
$\langle \tilde{\mathcal{C}} \rangle_m$	κ^2	κ^4/m	Gauss	approximates $\langle \mathcal{C} \rangle_m$ for large m

Notice that the digression on the shape of the PDFs can be omitted if the reader is interested only to the average and variance.

Next, we will analyze the properties of some useful estimators of \hat{S}_{yx} . Running an experiment, the logarithmic plot is comfortable because the average-to-deviation ratio is easily identified as the thickness of the track, independent of the vertical position. Yet, the logarithmic plot can only be used to display nonnegative quantities.

3.6.2 $\hat{S}_{yx} = |\langle S_{yx} \rangle_m|$

The main reason for us to spend attention with this estimator is that it is the default setting for cross-spectrum measurement in most FFT analyzers. Besides, it can be




Figure 3.8: Gaussian distribution of variance $\sigma^2 = 1/2$ and Rayleigh distribution generated by a pair of Gaussian variables of variance $\sigma^2 = 1/2$.

used in conjunction with $\arg \langle S_{yx} \rangle_m$ when the hypothesis that the delay of the two channels is not equal and useful information is contained in the argument, as it happens in radio-astronomy. $|\langle S_{yx} \rangle_m|$ is of course suitable to logarithmic plot because it can only take nonnegative values. The relevant objections against this estimator are

- There is no need to take in $\Im \{S_{yx}\}$, which contains half of the total background noise.
- The instrument background turns into relatively large estimation bias.

For large m , where $\langle \mathcal{C} \rangle_m$ tends to $\langle \tilde{\mathcal{C}} \rangle_m$, the estimator is expanded as

$$\begin{aligned} |\langle S_{yx} \rangle_m| &= \frac{1}{T} \sqrt{[\Re \{ \langle Y X^* \rangle_m \}]^2 + [\Im \{ \langle Y X^* \rangle_m \}]^2} \\ &= \frac{1}{T} \sqrt{[\langle \mathcal{A} \rangle_m + \langle \tilde{\mathcal{C}} \rangle_m]^2 + [\langle \mathcal{B} \rangle_m]^2}. \end{aligned}$$

The (not so) silly case of $\kappa = 0$

The analysis of this case tells us what happens when m is insufficient for the single-channel to be rejected, so that the displayed average spectrum is substantially the bias of the estimator. Since $c \leftrightarrow C = 0$, it holds that $\mathcal{C} = 0$. Letting

$$\langle \mathcal{Z} \rangle_m = \sqrt{[\langle \mathcal{A} \rangle_m]^2 + [\langle \mathcal{B} \rangle_m]^2}.$$

we notice that $\langle \mathcal{Z} \rangle_m$ is Rayleigh distributed with $2m$ degrees of freedom. Using Table 3.2, we find that

$$\mathbb{E}\{\langle \mathcal{Z} \rangle_m\} = \sqrt{\frac{\pi}{4m}} = \frac{0.886}{\sqrt{m}} \quad (\text{average})$$

$$\mathbb{V}\{\langle \mathcal{Z} \rangle_m\} = \frac{1}{m} \left(1 - \frac{\pi}{4}\right) = \frac{0.215}{m} \quad (\text{variance})$$




Figure 3.9: PDF of the estimator $\hat{S}_{yx} = \Re \{ \langle S_{yx} \rangle_m \}$.

Figure 3.8 compares the case $m = 1$ (Rayleigh distribution) to the Gaussian distribution associated with the best estimator (Section 3.6.3).

Interestingly, the deviation-to-average ratio, which also applies to $|\langle S_{yx} \rangle_m|$,

$$\frac{\text{dev}\{|\langle S_{yx} \rangle_m|\}}{\mathbb{E}\{|\langle S_{yx} \rangle_m|\}} = \sqrt{\frac{4}{\pi} - 1} = 0.523 \quad \frac{\text{dev}}{\mathbb{E}} \quad (3.18)$$

is independent of m . In logarithmic scale, the cross spectrum appears as a strip decreasing as $5 \log(m)$ dB, yet *of constant thickness* of approximately 5 dB (dev/avg). This is seen in the example of Fig. 3.6.

Large number of averaged realizations

The estimator converges to κ^2 , which is trivial, and for $\kappa \ll 1$ the deviation-to-average ratio is approximately $1/\sqrt{m}$. This issue is not further expanded here.

3.6.3 $\hat{S}_{yx} = \Re \{ \langle S_{yx} \rangle_m \}$

This is the best estimator to the extent that

- All the useful information is in $\Re \{ S_{yx} \} = \frac{1}{T}(\mathcal{A} + \mathcal{C})$.
- Since the instrument background is equally split in $\Re \{ S_{yx} \}$ and $\Im \{ S_{yx} \}$, discarding $\Im \{ S_{yx} \}$ results in 3 dB improvement of the SNR.
- for the same reason, the instrument background does not contribute to the bias.

The main drawback is that this estimator is not suitable to logarithmic plot because $\Re \{ \langle S_{yx} \rangle_m \}$ can take negative values, especially at small m . For large m we can approximate $\langle \mathcal{C} \rangle_m$ with $\langle \tilde{\mathcal{C}} \rangle_m$, which is Gaussian distributed. Letting

$$\langle \mathcal{Z} \rangle_m = \langle \mathcal{A} \rangle_m + \langle \tilde{\mathcal{C}} \rangle_m ,$$




Figure 3.10: PDF of the estimator $\hat{S}_{yx} = |\Re \{ \langle S_{yx} \rangle_m \}|$.

the PDF of $\langle \mathcal{Z} \rangle_m$ is Gaussian (Fig. 3.9). Using the results of Sec. 3.9.2, we find

$$\mathbb{E} \{ \langle \mathcal{Z} \rangle_m \} = \kappa^2 \quad (3.19)$$

$$\mathbb{V} \{ \langle \mathcal{Z} \rangle_m \} = \frac{1 + 2\kappa^2 + 2\kappa^4}{2m} \quad (3.20)$$

$$\text{dev} \{ \langle \mathcal{Z} \rangle_m \} = \sqrt{\frac{1 + 2\kappa^2 + 2\kappa^4}{2m}} \approx \frac{1 + \kappa^2}{\sqrt{2m}} \quad (3.21)$$

$$\frac{\text{dev} \{ \langle \mathcal{Z} \rangle_m \}}{\mathbb{E} \{ \langle \mathcal{Z} \rangle_m \}} = \frac{\sqrt{1 + 2\kappa^2 + 2\kappa^4}}{\kappa^2 \sqrt{2m}} \approx \frac{1 + \kappa^2}{\kappa^2 \sqrt{2m}} \quad (3.22)$$

$$P_N = \frac{1}{2} \text{erfc} \left(\frac{\kappa^2}{\sqrt{2} \sigma} \right) \quad (\mathbb{P} \{ \mathbf{x} < 0 \}, \text{ Sec. 3.9.2}) \quad (3.23)$$

$$P_P = 1 - \frac{1}{2} \text{erfc} \left(\frac{\kappa^2}{\sqrt{2} \sigma} \right) \quad (\mathbb{P} \{ \mathbf{x} > 0 \}, \text{ Sec. 3.9.2}) . \quad (3.24)$$

Accordingly, for $\kappa \ll 1$ a 0 dB SNR requires that $m = \frac{1}{2\kappa^4}$. If for example the DUT noise is 20 dB lower than the single-channel background, thus $\kappa = 0.1$, averaging on 5×10^3 spectra is necessary to get a SNR of 0 dB. On the other hand, if $\kappa \gg 1$ the deviation-to-average ratio converges to $1/\sqrt{2m}$, which is what we expect if the instrument background is negligible.

Precision vs. energy conservation

The term $\sqrt{2}$ in the denominator of (3.22) means that the SNR of the correlation system is 3 dB better than the single-channel system. In a physical system ruled by energy conservation this factor does not come for free because the DUT power is equally split into two channels. The conclusion is that the factor $\sqrt{2}$ in the SNR cancels with the $\sqrt{2}$ intrinsic loss of the power splitter. So, the basic *conservation laws* of thermodynamics (or information) are *not violated*.

3.6.4 $\hat{S}_{yx} = |\Re \{ \langle S_{yx} \rangle_m \}|$

The negative values of $\langle S_{yx} \rangle_m$ are folded up, so that \hat{S}_{yx} is always positive and can be plotted on a logarithmic axis. Approximating $\langle \mathcal{E} \rangle_m$ with $\langle \mathcal{E} \rangle_m$ for large m , the




Figure 3.11: PDF of the estimator obtained averaging the positive values of $\Re\{S_{yx}\}$.

estimator is expanded as

$$|\Re\{\langle S_{yx} \rangle_m\}| = \frac{1}{T} |\langle \mathcal{A} \rangle_m + \langle \tilde{\mathcal{C}} \rangle_m|$$

The PDF of $|\Re\{\langle S_{yx} \rangle_m\}|$ is obtained from the PDF of $|\Re\{\langle S_{yx} \rangle_m\}|$ already studied in Section 3.6.3 by folding² the negative-half-plane of the original PDF on the positive half plane. The result is shown in Fig. 3.10.

3.6.5 $\hat{S}_{yx} = \Re\{\langle S_{yx} \rangle_{m'}\}$, averaging on the positive values

Averaging m values of $\Re\{S_{yx}\}$, we expect $m' = m P_P$ positive values and $m - m' = m P_N$ negative values. This estimators consists of averaging on the m' positive values, discarding the negative values. As usual, we assume that for large m the term $\langle \mathcal{C} \rangle_m$ is approximated with $\langle \tilde{\mathcal{C}} \rangle_m$, so that its PDF is Gaussian. The PDF of this estimator is formed³ from the PDF of $\Re\{\langle S_{yx} \rangle_m\}$ after removing the negative-half-plane values and scaling up the result for the integral of the PDF to be equal to one. This is illustrated in Fig. 3.11.

3.6.6 Estimator $\hat{S}_{yx} = \langle \max(\Re\{S_{yx}\}, 0_+) \rangle_m$

Averaging $\Re\{S_{yx}\}$, the negative values are replaced with 0_+ . The reason for using 0_+ instead of just 0 is that $\lim_{x \rightarrow 0_+} \log(x)$ exists, while $\lim_{x \rightarrow 0} \log(x)$ does not. The notation “ 0_+ ” is a nerdish replacement for the “smallest positive floating-point number” available in the computer. This small number is equivalent to zero for all practical purposes, but never produces a floating-point error in the evaluation of the logarithm.

²A theorem states that follows. Let \mathbf{x} a random variable, $f(x)$ its PDF, and $\mathbf{y} = |\mathbf{x}|$ a function of \mathbf{x} . The PDF of \mathbf{y} is $g(y) = f(y)u(y) + f(-y)u(-y)$, where $u(y)$ is the Heaviside (step) function. Notice that the term $f(-y)u(-y)$ is the negative-half-plane ($y < 0$) side of $f(y)$ folded to the positive half plane.

³A theorem states that follows. Let $f(x)$ the PDF of a process, and $g(x)$ the PDF conditional to the event \mathbf{e} . The conditional PDF is obtained in two steps. First an auxiliary function $h(x)$ is obtained from $f(x)$ by selecting the sub-domain defined by \mathbf{e} . Second, the desired PDF is $g(x) = h(x) / \int_{-\infty}^{\infty} h(x) dx$. The first step generates $h(x)$ equal to $f(x)$, but taking away the portions not allowed by \mathbf{e} . The second step scales the function $h(x)$ up so that $\int_{-\infty}^{\infty} g(x) dx = 1$ (probability of all possible events), thus it is a valid PSD.





Figure 3.12: PDF of the estimator.

Figure 3.13: Comparison of the estimators based on $\Re\{S_{yx}\}$.

Since the negative values are replaced with zero, the PDF of this estimator (Fig. 3.12) derives from the PDF of $\Re\{\langle S_{yx} \rangle_m\}$ replacing the negative-half-plane side with a Dirac delta function.

3.6.7 Choice among the positive (biased) estimators

Having accepted that an estimator suitable to logarithmic plot is positive, thus inevitably biased, the best choice is the estimator that exhibits the lowest variance and the lowest bias. This criterion first excludes $|\langle S_{yx} \rangle_m|$ in favor of one of the estimators based on $\Re\{\langle S_{yx} \rangle_m\}$ because $\Im\{S_{yx}\}$ contains only the instrument background, which goes in both average (bias) and variance of $|\langle S_{yx} \rangle_m|$. Taking $\Im\{S_{yx}\}$ away, the estimator is necessarily based on $\Re\{\langle S_{yx} \rangle_m\}$.

Then, we search for a suitable low-bias estimator with the heuristic reasoning shown in Figure 3.13.

It is shown in Sec. 3.6.3 that for large m the PDF of $\Re\{\langle S_{yx} \rangle_m\}$ is a Gaussian distribution with mean value κ^2 and variance $\sigma^2 = \frac{1+2\kappa^2+2\kappa^4}{2m}$. The probability of the events $\Re\{\langle S_{yx} \rangle_m\} < 0$ is represented in Fig. 3.9 as the grey area on the left-hand half-plane. These events have probability P_N . Using the results of Section 3.9.2, the average of these negative events is

$$\mu_N = \int_{-\infty}^{\infty} x f_N(x) dx = \mu - \frac{1}{\frac{1}{2}\text{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right)} \frac{\sigma}{\sqrt{2\pi \exp(\mu^2/\sigma^2)}} \quad (\text{Eq. (3.40)}) .$$

The estimator is made positive by moving the area P_N from the left-hand half-plane to the right-hand half-plane. The bias depends on the shape taken by this area, and ultimately on the average associated to this shifted P_N . By inspection on Fig. 3.13 we notice that

Section 3.6.5. $\hat{S}_{yx} = \Re\{\langle S_{yx} \rangle_{m'}\}$ makes use only of the positive values, the negative values are discarded. The PSD area associated to P_N has the same shape of the right-hand side of the PSD. We denote the average of this shape with μ_1 .

Section 3.6.4. $\hat{S}_{yx} = |\Re\{\langle S_{yx} \rangle_m\}|$. The shadowed area associated to P_N is flipped from the negative half-plane to the positive half-plane. The average is $\mu_2 = -\mu_N$.

Section 3.6.6. $\hat{S}_{yx} = \Re\{\langle \max(S_{yx}, 0_+) \rangle_m\}$. The shadowed area associated to P_N collapses into a Dirac delta function. The average is $\mu_3 = 0$.

From the graphical construction of Fig. 3.13, it is evident that

$$\mu_1 > \mu_2 > \mu_3 .$$

The obvious conclusion is that the preferred estimator is

$$\hat{S}_{yx} = \Re\{\langle \max(S_{yx}, 0_+) \rangle_m\} \quad (\text{Preferred, Sec. 3.6.6}) .$$

It is worth pointing out that the naïf approach of just *discarding the negative values* before averaging (Sec. 3.6.5) turns out to be the *worst choice* among the estimators we analyzed.

3.6.8 The use of $\Im\{\langle S_{yx} \rangle_m\}$

It has been shown in Sec. 3.6 (Eq. (3.16)) that all the DUT signal goes into $\Re\{S_{yx}\}$, and that $\Re\{S_{yx}\}$ contains only the instrument background. More precisely, (3.16) is rewritten as

$$S_{yx} = \frac{1}{T} \mathbb{E} \{ \mathcal{A} + i\mathcal{B} + \mathcal{C} \} \quad (\text{Eq. (3.17)})$$

$$\Re\{S_{yx}\} = \frac{1}{T} \mathbb{E} \{ \mathcal{A} + \mathcal{C} \} \quad \text{and} \quad \Im\{S_{yx}\} = \frac{1}{T} \mathbb{E} \{ \mathcal{B} \}$$

where \mathcal{A} and \mathcal{B} come from the background have equal statistics, and \mathcal{C} comes from the DUT spectrum. Therefore

- $\Im\{\langle S_{yx} \rangle_m\}$ is a good estimator of the background
- the contrast $\Re\{\langle S_{yx} \rangle_m\} - \Im\{\langle S_{yx} \rangle_m\}$ is a good indicator of the averaging convergence to S_{cc} .




Figure 3.14: Effect of the finite duration of the measurement on the spectrum.

3.7 Statistical independence on the frequency axis

As a relevant property of white Gaussian noise, the Fourier transform is also Gaussian with all values on the frequency axis statistically-independent. This property is taken as a good representation of the reality even in the case of discrete spectra measured on a finite measurement time T , and used extensively in this report. Yet, in a strictly mathematical sense time-domain truncation breaks the hypothesis of statistical independence in the frequency domain. This happens because time truncation is equivalent to a multiplication by a rectangular pulse, which maps into a convolution by a $\text{sinc}()$ function in the frequency domain. This concept is shown in Fig. 3.14, and expanded as follows

$$\begin{aligned}
 x(t) &\Rightarrow x_T(t) = x(t) \Pi(t/T) \\
 X(f) &\Rightarrow X_T(f) = x(t) * T \frac{\sin(\pi T f)}{\pi T f}
 \end{aligned}$$

where

$$\Pi(t) = \begin{cases} 1 & -1/2 < t < 1/2 \\ 0 & \text{elsewhere} \end{cases} \quad \Leftrightarrow \quad \text{sinc}(f) = \frac{\sin(\pi f)}{\pi f} .$$

The consequences are the following.

- The side-lobes of $T\text{sinc}(Tf)$ cause energy leakage, thus a small correlation on the frequency axis.

- Accuracy is reduced because each point collects energy from other frequencies. This may show up in the presence of high peaks (50–60Hz, for example) or high roll-off bumps.
- One should question whether the number of degrees of freedom is reduced.

The truncation function is called “window” on the front panel of analyzers, and sometimes “taper” in textbooks about spectral analysis. Reduced frequency leakage is obtained by a different choice of the truncation function, like the Bartlett (triangular), Hanning (cosine) or Parzen (cubic) window.

3.8 Applications and experimental techniques

3.8.1 PM noise

The first application to frequency metrology was the measurement of Hydrogen masers [?] in the early sixties. Then, the method was used for the measurement of phase noise [?] in the seventies, but it found some popularity only in the nineties, when dual-channel FFT analyzers started to be available.

Figure 3.15 shows some of the most popular schemes for the measurement of phase noise. The mixer is a saturated phase-to-voltage converter in Fig. 3.15 A-C, and a synchronous down-converter in Fig. 3.15 D. In all cases correlation is used to reject the noise of the two mixers. The background noise turns out to be limited by the thermal homogeneity, instead of the absolute temperature referred to the carrier power. This property was understood only after working on the scheme D [?]. At that time, the other schemes were already known.

The scheme A [?] is suitable to the measurement of low-noise two-port devices, mainly passive devices showing small group delay, so that the noise of the reference oscillator can be rejected.

The scheme B consists of two separate PLLs that measure separately the oscillator under test. Correlation rejects the noise of the two reference oscillators. In this way, it is possible to measure an oscillator by comparing it to a pair of synthesizers, even if the noise of the synthesizers is higher than that of the oscillator. This fact is relevant to the development of oscillator technology, when manufacturing makes it difficult to have the oscillator at the round frequency of the available standards, and also difficult to build two prototypes at the same frequency.

The scheme C derives from A after introducing a delay in the arms [?]. It can be implemented using either a pair of resonators or a pair a delay lines. The use of the optical-fiber delay line is the most promising solution because the delay line can be adapted to the arbitrary frequency of the oscillator under test, while a resonator can not [?]. Correlation removes the fluctuations of the delay line [?, ?].

The scheme D is based on a bridge that nulls the carrier before amplification and synchronous detection of the noise sidebands. This scheme derives from the pioneering work of Sann [?]. At that time, the mixer was used to down convert the fluctuation of the null at the output of a magic Tee. Amplification of the noise sideband [?] and correlation [?] were introduced afterwards.

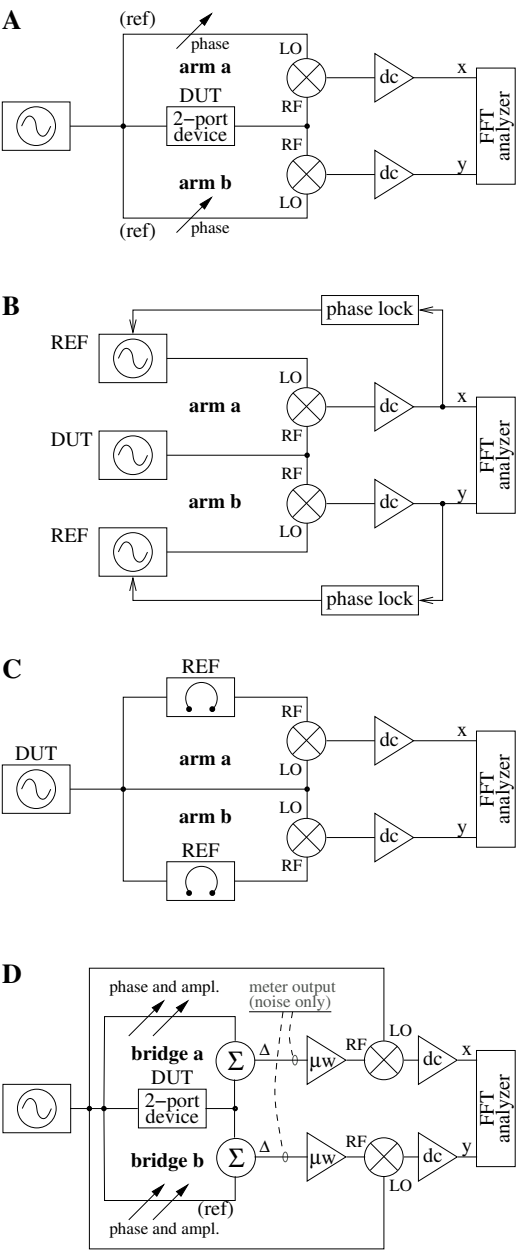


Figure 3.15: Basics schemes for the measurement of phase noise.

With modern RF/microwave components, isolation between the two channels may not be a serious problem. The hardware sensitivity is limited environmental effects, like temperature fluctuations and low-frequency magnetic fields, and by the AM noise.


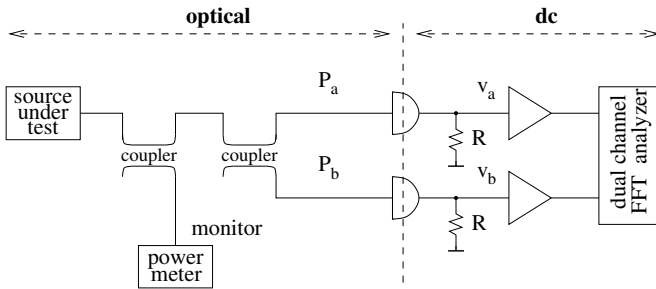
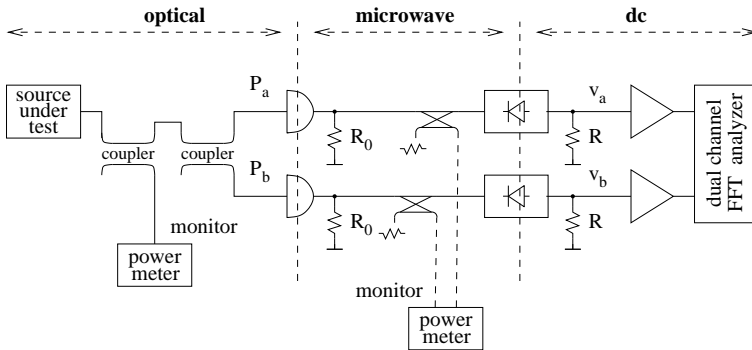
A: amplitude noise of a RF/microwave source**B: relative intensity noise (RIN) of a laser****C: amplitude noise of a photonic RF/microwave source**

Figure 3.16: Basics schemes for the measurement of amplitude noise (from [?]).

The latter is taken in through the sensitivity of the mixer offset to the input power. Only partial solutions are available [?].




Figure 3.17: Example of cross spectrum measurement (amplitude noise of an oven-controlled quartz oscillator), taken from [?].

3.8.2 AM noise

Figure 3.16 shows some schemes for the cross spectrum measurement of AM noise, taken from [?].

In Fig. 3.16 A, two Schottky-diode or tunnel-diode passive power-detectors are used to measure simultaneously the power fluctuations of the source under test. Isolation between channels is guaranteed by the isolation of the power splitter (18–20 dB) and by the fact that the power detectors do not send noise back to the input. Correlation enables the rejection the single-channel noise.

As an example, Fig. 3.17 shows the measurement of a quartz oscillator. Converting the $1/f$ noise into stability of the fractional amplitude α , we get $\sigma_\alpha(\tau) = 4.3 \times 10^{-7}$ (Allan deviation, constant vs. the measurement time τ). This oscillator exhibits the lowest AM noise measured in our laboratory. The single-channel noise rejection achieved by correlation and averaging is more than 10 dB.

Figure 3.16 B is the obvious adaptation of the scheme A to the measurement of the laser relative intensity noise (RIN). We start using it routinely.

The scheme of Fig. 3.16 C, presently under study, is intended for the measurement of the microwave AM noise on the modulated light beam at the output of new generation of opto-electronic oscillators based on optical fibers [?], or based on whispering-gallery optical resonators.

Single-channel vs. dual-channel measurements

In the measurement of PM noise it is more or less possible to test the background of a single-channel instrument by removing the DUT. This happens because we can always get the two phase-detector from a single oscillator, which is the phase reference.⁴ The correlation schemes are more complex than the single-channel counterparts, and some-

⁴This statement of course applies only to the background noise of the instrument. When the instrument is used to measure an oscillator we need a reference oscillator, the noise of which must be validated separately.

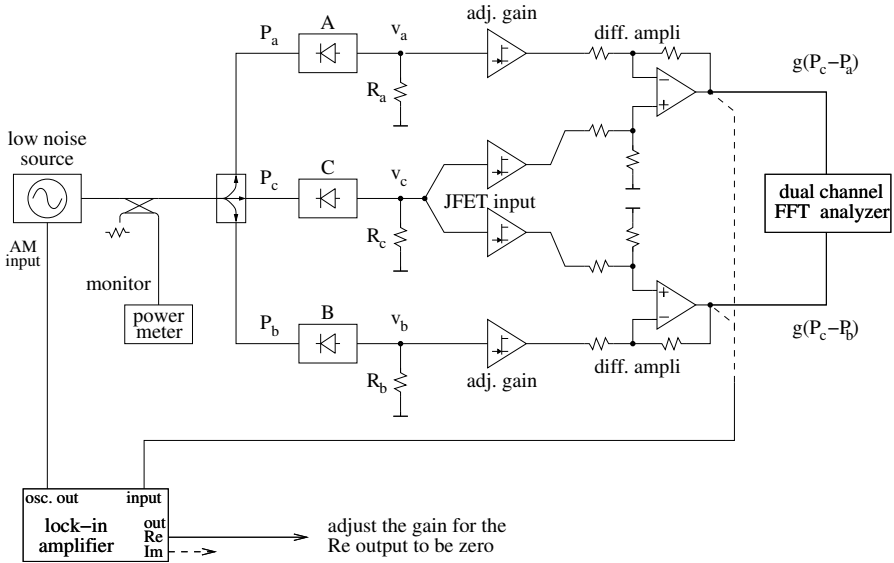


Figure 3.18: Measurement of the background noise of a power detector.

times difficult to operate. Obviously, the experimentalist prefers the single-channel measurements and uses the correlation schemes only when the sensitivity of the former is insufficient.

Conversely, the measurement of AM noise relies upon the power detector, which does not work without the source. Thus we cannot remove the device under test, and of course we cannot assess the single-channel background noise of the instrument in this way. One can object that even in the case of PM noise we can not measure an oscillator in single-channel mode if we do not have a low-noise reference oscillator. The difference is that in the case of PM noise we can at least validate the instrument, while in the case of AM noise we can not.

Another difference between AM and PM is that the phase detector is always more or less sensitive to AM noise [?], while the amplitude detector is not sensitive to phase noise. In correlation systems, this fact makes the channel separation simple to achieve and to test.

The conclusion is that the cross-spectrum measurement is inherently simpler with AM noise than with PM noise.

3.8.3 Other applications

Tracking back through the literature, the first use of the cross-spectrum was for the determination of the angular size of stellar radio sources [?]. In the case of a signal coming through two antennas separated by an appropriate baseline, the latter introduces a delay depending on the source direction in space. Hence the useful signal S_{cc} cannot be real. Instead, the angle $\arctan \Im/\Re$ gives information on the source direction. The very-large-baseline interferometry (VLBI) can be seen as a generalization

of this method.

When the same method was applied to the intensity interferometer [?,?], an anti-correlation effect was discovered, due to the discrete nature of light. This phenomenon, known as Hanbury Brown – Twiss effect (HBT effect), was later observed also in microwave signals in photonic regime [?], i.e., with $h\nu > kT$.

The correlation method finds another obvious application in radiometry [?], and of course in Johnson thermometry, which is often considered a branch of radiometry.

Since the cross-spectrum enables to compare the PSD of two noise sources, it can be used to measure a temperature by comparing thermal noise to a reference shot noise. The latter is in turn measured as a dc value by exploiting the property of Poisson processes that the variance can be calculated from the average. In a Tunnel junction, theory predicts the amount of shot and thermal noise. This fact can be exploited for precision thermometry [?], and ultimately to redefine the temperature in terms of fundamental constants.

The measurement of the low $1/f$ voltage fluctuations is an important diagnostic tool in semiconductor technology. The field-effect transistors are suitable to this task because of the low bias current at the input. In fact, the bias current flowing into the sample turns into a fully correlated voltage through the Ohm law. Additionally, the electrode capacitance may limit the instrument sensitivity. The reader can refer to [?] for a detailed treatise.

In metallurgy, the cross spectrum method has been used for the measurement of electromigration in thin metal films through the $1/f$ fluctuation of the conductor resistance. This is relevant in microprocessor technology because the high current density in metal connexions can limit the life of the component and make it unreliable. For this reason, Aluminum is no longer used. The high sensitivity is based on the idea that with white Gaussian noise X' and X'' (real and imaginary part) are statistically independent. Synchronously detecting the signal with two orthogonal references, it is therefore possible to reject the amplifier noise even if a single amplifier is shared by the two channel [?]. Adapting this idea to RF and microwaves is straightforward [?]. Unfortunately, we still have no application for this.

3.9 AppendixMathematical background

3.9.1 Random variables and density functions

Let \mathbf{x} a random variable and x a variable. Denoting with $\mathbb{P}\{\mathbf{e}\}$ the probability of the event \mathbf{e} , two relevant probability functions are associated with \mathbf{x} and x , namely the cumulative density function $F(x)$ and the probability density function $f(x)$. They are defined as

$$F(x) = \mathbb{P}\{\mathbf{x} < x\} \quad (\text{cumulative density function, or CDF}) \quad (3.25)$$

$$f(x) dx = \mathbb{P}\{x < \mathbf{x} < x + dx\} \quad (\text{probability density function, or PDF}) . \quad (3.26)$$




Figure 3.19: Gaussian (normal) PDF.

CDF and PDF are related by

$$F(x) = \int_{-\infty}^x f(x') dx' . \quad (3.27)$$

The probability that \mathbf{x} is in the interval $[a, b]$ is

$$\mathbb{P}\{a < \mathbf{x} < b\} = F(b) - F(a) = \int_a^b f(x) dx . \quad (3.28)$$

The probability that \mathbf{x} takes any value is equal to one, thus

$$F(\infty) = 1 \quad \text{and} \quad \int_{-\infty}^{\infty} f(x) dx = 1 . \quad (3.29)$$

The average and the variance of the random variable \mathbf{x} are

$$\mathbb{E}\{\mathbf{x}\} = \int_{-\infty}^{\infty} x f(x) dx \quad \mathbb{E}\{\mathbf{x}\}, \text{ (average)} \quad (3.30)$$

$$\mathbb{V}\{\mathbf{x}\} = \mathbb{E}\{|\mathbf{x} - \mathbb{E}\{\mathbf{x}\}|^2\} = \int_{-\infty}^{\infty} (x - \mathbb{E}\{\mathbf{x}\})^2 f(x) dx \quad \text{(variance)} \quad (3.31)$$

3.9.2 Gaussian (normal) distribution (Fig. 3.19)

The Gaussian (normal) distribution has the following main properties

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right] \quad (\text{Gaussian PDF}) \quad (3.32)$$

$$\mathbb{E}\{\mathbf{x}\} = \mu \quad (\text{average}) \quad (3.33)$$

$$\mathbb{V}\{\mathbf{x}\} = \sigma^2 \quad (\text{variance}) \quad (3.34)$$

$$P_N = \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right) \quad (\mathbb{P}\{\mathbf{x} < 0\}) \quad (3.35)$$

$$P_P = 1 - \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right) \quad (\mathbb{P}\{\mathbf{x} > 0\}) \quad (3.36)$$

A new PDF is associated to the positive events

$$f_P(x) = \frac{1}{P_P} f(x) \mathbf{u}(x) \quad (3.37)$$

$$\mu_P = \int_{-\infty}^{\infty} x f_P(x) dx = \mu + \frac{1}{1 - \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right)} \frac{\sigma}{\sqrt{2\pi \exp(\mu^2/\sigma^2)}} \quad (3.38)$$

Similarly, another PDF is associated to the negative events

$$f_N(x) = \frac{1}{P_N} f(x) \mathbf{u}(-x) \quad (3.39)$$

$$\mu_N = \int_{-\infty}^{\infty} x f_N(x) dx = \mu - \frac{1}{\frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right)} \frac{\sigma}{\sqrt{2\pi \exp(\mu^2/\sigma^2)}} \quad (3.40)$$

The following integrals related to the Gaussian PDF are useful

$$\int_{-\infty}^0 f(x) dx = \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right) \quad (3.41)$$

$$\int_0^{\infty} f(x) dx = 1 - \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right) \quad (3.42)$$

$$\int_{-\infty}^0 x f(x) dx = \mu \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right) - \frac{\sigma}{\sqrt{2\pi \exp(\mu^2/\sigma^2)}} \quad (3.43)$$

$$\int_0^{\infty} x f(x) dx = \mu \left[1 - \frac{1}{2} \operatorname{erfc}\left(\frac{\mu}{\sqrt{2}\sigma}\right)\right] + \frac{\sigma}{\sqrt{2\pi \exp(\mu^2/\sigma^2)}} \quad (3.44)$$

Sum of zero-mean Gaussian variables

Let $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ two random functions with Gaussian distribution, zero mean and variance σ_1^2 and σ_2^2 . The sum $\mathbf{x}(t) = \mathbf{x}_1(t) + \mathbf{x}_2(t)$ is a random function with Gaussian distribution, zero mean and variance $\sigma^2 = \sigma_1^2 + \sigma_2^2$.

Sum of a nonzero-mean and a zero-mean Gaussian variable

Let $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ two random functions with Gaussian distribution, mean $\mu_1 \neq 0$ and $\mu_2 = 0$, and variance σ_1^2 and σ_2^2 . The sum $\mathbf{x}(t) = \mathbf{x}_1(t) + \mathbf{x}_2(t)$ is a random function with Gaussian distribution, mean $\mu = \mu_1$ and variance $\sigma^2 = \sigma_1^2 + \sigma_2^2$.

Product of zero-mean Gaussian variables

Let $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ two random functions with Gaussian distribution, zero mean and variance σ_1^2 and σ_2^2 . The product $\mathbf{x} = \mathbf{x}_1(t) \mathbf{x}_2(t)$ is a random function with *Bessel-K* distribution [?] (aka, *product-normal* distribution), zero mean and variance $\sigma^2 = \sigma_1^2 \sigma_2^2$. The Bessel-*K* distribution is defined as

$$\frac{1}{\pi \sigma_1 \sigma_2} K_0 \left(\frac{|x|}{\sigma_1 \sigma_2} \right), \quad (3.45)$$

where

$$K_0(x) = \int_0^\infty \cos(x \sinh t) dt = \int_0^\infty \frac{\cos(xt)}{\sqrt{t^2 + 1}} dt \quad (3.46)$$

is the Bessel function of the second kind, and the factor $\frac{1}{\pi \sigma_1 \sigma_2}$ is necessary to make the integral equal to one.

Average product of zero-mean Gaussian variables

Averaging on m realizations on the product $\mathbf{x} = \mathbf{x}_1(t) \mathbf{x}_2(t)$ as above, the distribution is the convolution of m equal Bessel-*K* distributions. Thanks to the *central limit* theorem, for large m the distribution converges to a Gaussian distribution with zero mean and variance $\sigma^2 = \sigma_1^2 \sigma_2^2$.

Fourier transform of a Gaussian variable

Let $\mathbf{x}(t)$ a random process with Gaussian distribution and white spectrum, and $x(t)$ a realization. The Fourier transform $X(f) = X'(f) + \imath X''(f)$ is a random process with white spectrum and zero-mean Gaussian distribution. This means that

1. At any frequency, the real part $X'(f)$ and the imaginary part $X''(f)$ are random variables statistically independent with equal variance.
2. Given two frequencies f_1 and f_2 (or two separate frequency intervals), $X(f_1)$ and $X(f_2)$ are statistically independent.

Interestingly, $|X| = \sqrt{(X')^2 + (X'')^2}$ has Rayleigh distribution, and $|X|^2 = (X')^2 + (X'')^2$ has χ^2 distribution with 2 degrees of freedom.

Discrete zero-mean Gaussian-distributed white noise

It is often convenient to use the discrete Fourier transform and spectra. Thus we refer to

$$\begin{aligned} X(f) &\Rightarrow X_{ij} = X'_{ij} + \imath X''_{ij} \\ S(f) &\Rightarrow S_{ij} = \frac{1}{T} (X'_{ij}{}^2 + X''_{ij}{}^2) \end{aligned}$$

where the subscript i denotes the i -th realization and the subscript j denotes the discrete frequency. The following properties hold for zero-mean white noise with Gaussian distribution.

1. X_{ij} is zero-mean Gaussian distributed. Thus X'_{ij} and X''_{ij} are zero-mean Gaussian processes.
2. Different frequency.
 - X_{ij} and X_{ik} , $j \neq k$, are statistically independent.
 - $\mathbb{V}\{X_{ij}\} = \mathbb{V}\{X_{ik}\}$ (energy equipartition).
3. Real and imaginary part.
 - X'_{ij} and X''_{ij} are statistically independent.
 - $\mathbb{E}\{X'_{ij}\} = 0$, and $\mathbb{E}\{X''_{ij}\} = 0$ (zero mean).
 - $\mathbb{V}\{X'_{ij}\} = \mathbb{V}\{X''_{ij}\} = \frac{1}{2}\mathbb{V}\{X_{ij}\}$ (energy equipartition).
4. Absolute square value $|X_{ij}|^2 = |X'_{ij}|^2 + |X''_{ij}|^2$. Letting $\mathbb{V}\{X_{ij}\} = \sigma^2$,
 - $|X_{ij}|^2$ has χ^2 distribution with two degrees of freedom.
 - $\mathbb{E}\{|X_{ij}|^2\} = \sigma^2$ (average).
 - $\mathbb{V}\{|X_{ij}|^2\} = \sigma^4$ (variance).
5. Sum of two independent processes, $\mathbf{y} = \mathbf{x}_1 + \mathbf{x}_2 \leftrightarrow Y = X_1 + X_2$.
 - Y_{ij} is Gaussian distributed
 - $\mathbb{V}\{Y_{ij}\} = \mathbb{V}\{X_{1ij}\} + \mathbb{V}\{X_{2ij}\}$.
6. Product of two independent processes, $\mathbf{y} = \mathbf{x}_1 \mathbf{x}_2 \leftrightarrow Y = X_1 * X_2$.
 - Y_{ij} is normal-product distributed
 - $\mathbb{V}\{Y_{ij}\} = \mathbb{V}\{X_{1ij}\} \mathbb{V}\{X_{2ij}\}$.

3.9.3 Chi-square distribution (Table 3.1)

Let $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\nu$ a set of normal-distributed random variables with zero mean and variance equal one, and

$$\chi^2 = \sum_{i=1}^{\nu} \mathbf{x}_i^2 \quad (3.47)$$

a new function called ‘chi-square’ distribution with ν degrees of freedom. The probability functions associated to $\mathbf{x} = \chi^2$ and the relevant parameters are

$$f(x) = \frac{x^{\frac{1}{2}\nu-1} e^{-\frac{1}{2}x}}{\Gamma(\frac{1}{2}\nu) 2^{\frac{1}{2}\nu}} \quad x \geq 0 \quad (\text{chi-square PDF}) \quad (3.48)$$

$$F(x) = 1 - \frac{\Gamma(\frac{1}{2}\nu, \frac{1}{2}x)}{\Gamma(\frac{1}{2}\nu)} = \frac{\gamma(\frac{1}{2}\nu, \frac{1}{2}x)}{\Gamma(\frac{1}{2}\nu)} \quad (\text{chi-square CDF}) \quad (3.49)$$

$$\mathbb{E}\{\mathbf{x}\} = \nu \quad (\text{average}) \quad (3.50)$$

$$\mathbb{E}\{\mathbf{x}^2\} = \nu(\nu + 2) \quad (\text{2nd moment}) \quad (3.51)$$

$$\mathbb{E}\{|\mathbf{x} - \mathbb{E}\{\mathbf{x}\}|^2\} = 2\nu \quad (\text{variance}) . \quad (3.52)$$

It follows immediately from the definition of χ^2 that the sum of n random variables with χ^2 distribution and ν_j degrees of freedom is χ^2 distributed

$$\chi^2 = \sum_{j=1}^n \chi_j^2, \quad \nu = \sum_{j=1}^n \nu_j .$$

In the general case, the variance of $\mathbf{x}_1 \dots \mathbf{x}_\nu$ is $\sigma^2 \neq 1$. This is solved with the transformation $\mathbf{x} \rightarrow \mathbf{x}/\sigma^2$. Thus $f(x) = \frac{1}{\sigma^2} [f(x)]_{\text{var}=1}$, and⁵

$$f(x) = \frac{x^{\frac{1}{2}\nu-1} e^{-\frac{1}{2}\frac{x}{\sigma^2}}}{\sigma^\nu \Gamma(\frac{1}{2}\nu) 2^{\frac{1}{2}\nu}} \quad x \geq 0 \quad (\text{chi-square PDF}) \quad (3.53)$$

$$\mathbb{E}\{\mathbf{x}\} = \sigma^2 \nu \quad (\text{average}) \quad (3.54)$$

$$\mathbb{E}\{\mathbf{x}^2\} = \sigma^4 \nu(\nu + 2) \quad (\text{2nd moment}) \quad (3.55)$$

$$\mathbb{E}\{|\mathbf{x} - \mathbb{E}\{\mathbf{x}\}|^2\} = 2\sigma^4 \nu \quad (\text{variance}) . \quad (3.56)$$

3.9.4 Rayleigh distribution

Let \mathbf{x}_1 and \mathbf{x}_2 , two independent random functions with Gaussian distribution, zero mean and equal variance σ , and

$$\mathbf{x} = \sqrt{\mathbf{x}_1^2 + \mathbf{x}_2^2} \quad (3.57)$$

⁵Cramér p. 236

Table 3.1: Some properties of the χ^2 distribution.

distribution	χ^2	scaled χ^2	scaled χ^2 , average of m real \mathbf{x}_i^2	scaled χ^2 , average of m complex $\mathbf{x}_i^2 = \mathbf{x}_i'^2 + \mathbf{x}_i''^2$
rule	$\mathbf{x} = \chi^2 = \sum_{i=1}^{\nu} \mathbf{x}_i^2$ $\mathbb{E}\{\mathbf{x}_i\} = 0$ $\mathbb{V}\{\mathbf{x}_i\} = 1$	$\mathbf{x} = \chi^2 = \sum_{i=1}^{\nu} \mathbf{x}_i^2$ $\mathbb{E}\{\mathbf{x}_i\} = 0$ $\mathbb{V}\{\mathbf{x}_i\} = \sigma^2$	$\mathbf{x} = \frac{1}{m} \sum_{i=1}^{\nu} \mathbf{x}_i^2$ $\mathbb{E}\{\mathbf{x}_i\} = 0$ $\mathbb{V}\{\mathbf{x}_i\} = \sigma^2$	$\mathbf{x} = \frac{1}{m} \sum_{i=1}^{\nu} (\mathbf{x}_i'^2 + \mathbf{x}_i''^2)$ $\mathbb{E}\{\mathbf{x}_i'\} = \mathbb{E}\{\mathbf{x}_i''\} = 0$ $\mathbb{V}\{\mathbf{x}_i'\} = \mathbb{V}\{\mathbf{x}_i''\} = \sigma^2$
transformation	none $f(x) = k_{\nu}(x)$	$f(x) = \frac{1}{\sigma^2} k_{\nu}\left(\frac{x}{\sigma^2}\right)$	$f(x) = \frac{m}{\sigma^2} k_m\left(\frac{mx}{\sigma^2}\right)$	$f(x) = \frac{2m}{\sigma^2} k_{2m}\left(\frac{2mx}{\sigma^2}\right)$
probability density function	$f(x) = \frac{x^{\frac{1}{2}\nu-1} e^{-\frac{1}{2}x}}{\Gamma(\frac{1}{2}\nu) 2^{\frac{1}{2}\nu}}$	$f(x) = \frac{x^{\frac{1}{2}\nu-1} e^{-\frac{x}{2\sigma^2}}}{\sigma^{\nu} \Gamma(\frac{1}{2}\nu) 2^{\frac{1}{2}\nu}}$	$f(x) = \frac{m^{\frac{1}{2}m} x^{\frac{1}{2}m-1} e^{-\frac{mx}{2\sigma^2}}}{\sigma^m \Gamma(\frac{1}{2}m) 2^{\frac{1}{2}m}}$	$f(x) = \frac{m^m x^{\frac{1}{2}m-1} e^{-\frac{mx}{\sigma^2}}}{\sigma^{2m} \Gamma(m)}$
average $\mathbb{E}\{\mathbf{x}\}$	ν	$\nu\sigma^2$	σ^2	σ^2
variance $\mathbb{E}\{ \mathbf{x} - \mathbb{E}\{\mathbf{x}\} ^2\}$	2ν	$2\nu\sigma^4$	$2\sigma^4/m$	σ^4/m

a new random function. This function has Rayleigh probability density function⁶

$$f(x) = \frac{x}{\sigma^2} \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad y > 0 \quad (\text{Rayleigh PDF}) \quad (3.58)$$

Table 3.2: Relevant values for the Rayleigh distribution.

Rayleigh distribution with $\sigma^2 = 1/2$	
quantity with $\sigma^2 = 1/2$	value [10 log(), dB]
average = $\sqrt{\frac{\pi}{4}}$	0.886 [-0.525]
deviation = $\sqrt{1 - \frac{\pi}{4}}$	0.463 [-3.34]
$\frac{\text{dev}}{\text{avg}} = \sqrt{\frac{4}{\pi} - 1}$	0.523 [-2.8]
$\frac{\text{avg} + \text{dev}}{\text{avg}} = 1 + \sqrt{\frac{4}{\pi} - 1}$	1.523 [+1.83]
$\frac{\text{avg} - \text{dev}}{\text{avg}} = 1 - \sqrt{\frac{4}{\pi} - 1}$	0.477 [-3.21]

The functions $\mathbf{x}_1(t)$ and $\mathbf{x}_2(t)$ can be interpreted as the random amplitude of two orthogonal vectors, or the real and imaginary part of a complex random function. Following this interpretation, $\mathbf{x}(t)$ is the absolute value of the vector sum. Table 3.2 reports some useful numerical values related to the $\sigma^2 = 1/2$ Rayleigh distribution.

A case of interest in averaged measurement is $\sigma^2 = 1/2m$, which yields

$$\mathbb{E}\{\mathbf{x}\} = \sqrt{\frac{\pi}{4m}} = \frac{0.886}{\sqrt{m}} \quad (\text{average}) \quad (3.62)$$

$$\mathbb{E}\{\mathbf{x}^2\} = \frac{1}{m} \quad (2\text{nd moment}) \quad (3.63)$$

$$\mathbb{V}\{\mathbf{x}\} = \mathbb{E}\{|\mathbf{x} - \mathbb{E}\{\mathbf{x}\}|^2\} = \left(1 - \frac{\pi}{4}\right) \frac{1}{m} = \frac{0.215}{m} \quad (\text{variance}) \quad (3.64)$$

$$\sqrt{\mathbb{V}\{\mathbf{x}\}} = \sqrt{\frac{1}{m} \left(1 - \frac{\pi}{4}\right)} = \frac{0.463}{\sqrt{m}} \quad (\text{deviation}) \quad (3.65)$$

$$\frac{\sqrt{\mathbb{V}\{\mathbf{x}\}}}{\mathbb{E}\{\mathbf{x}\}} = \sqrt{\frac{4}{\pi} - 1} = 0.523 \quad (\text{independent of } m) \quad (\text{dev/avg}) . \quad (3.66)$$

3.10 Windowing

3.11 Overlapping

Chapter 4

Allan variances

The classical variance¹

$$\sigma^2 = \frac{1}{n-1} \sum_{i=1}^n [x_i - \mu]^2, \quad (4.1)$$

where $\mu = \frac{1}{n} \sum_{i=1}^n x_i$ is the average, fails at describing time divergent processes because (i) it depends on the averaging time used to take the samples x_i , and (ii) it depends on the number n of samples. Try this yourself, feeding $x_i = 1.0001, 1.0002, 1.0003 \dots$, with $n = 2, 4, 8 \dots$ in (4.1). A solution consists of introducing the averaging time as a parameter, denoted with τ , and to set $n = 2$. This is the minimum n which gives a valid σ^2 . Welcome to the Allan variance.

4.0.1 Definition and Evaluation (Region 1.3)

The two-sample (Allan) variance AVAR of the quantity y is defined as

$$\text{2-sample variance: } \sigma_y^2(\tau) = \mathbb{E} \left\{ \frac{1}{2} [\bar{y}_2 - \bar{y}_1]^2 \right\}, \quad (4.2)$$

where $\mathbb{E}\{ \}$ is the mathematical expectation, and the averages \bar{y}_1 and \bar{y}_2 are taken over contiguous time slots of duration τ . The quantity $\frac{1}{2} [\bar{y}_2 - \bar{y}_1]^2$ is the classical

¹This is a simplified notation. More precisely, (4.1) describes an estimator, thus it should be written as $\widehat{\sigma^2} = \dots [x_i - \hat{\mu}]^2$, where $\hat{\mu} = \dots$

variance evaluated with two samples. This is immediately seen by replacing $n = 2$ in (4.1), then $\mu = (x_1 + x_2)/2$; finally, we identify the generic x_i with the average fractional frequency fluctuation \bar{y}_i . The quantity ADEV, a deviation, the square root of AVAR—and similarly MDEV, PDEV etc. defined later—can be seen as an estimator of the uncertainty of the quantity y , accumulated in the time τ after reset or calibration.

In experiments, $\mathbb{E}\{\}$ is replaced with the estimator

$$\widehat{\sigma_y^2}(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} [\bar{y}_{k+1} - \bar{y}_k]^2, \quad (4.3)$$

which is the average on $M-1$ realizations of $\bar{y}_2 - \bar{y}_1$, thus it requires M contiguous measures of \bar{y} .

At first reading, one can take (4.3) as the formula to evaluate AVAR, ignoring the ‘hat.’ Some Authors use (4.3) as the definition of AVAR. That said, keeping a clear difference between $\sigma_y^2(\tau)$ and its estimate $\widehat{\sigma_y^2}(\tau)$ is important in theoretical analysis and in the evaluation of the uncertainty.

It is often convenient to rely on time measurements x_k , using $y_k = (x_{k+1} - x_k)/\tau$. Accordingly, (4.3) rewrites as

$$\widehat{\sigma_y^2}(\tau) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} \left[\frac{x_{k+2} - 2x_{k+1} + x_k}{\tau} \right]^2, \quad (4.4)$$

which requires $M+1$ measures of $x(t)$ spaced by τ .

4.0.2 Spectral Response (Region 1.5)

The Allan variance can be calculated from the spectrum using

$$\sigma_y^2(\tau) = \int_0^\infty |H_A(f; \tau)|^2 S_y(f) df, \quad (4.5)$$

where the transfer function

$$|H_A(f; \tau)|^2 = 2 \frac{\sin^4(\pi\tau f)}{(\pi\tau f)} \quad (4.6)$$

is similar to an octave bandpass filter centered at $f \simeq 0.45/\tau$. Unfortunately, such filter suffers from significant side lobes (inset in the plot of Region 1.5). Notice that $\sigma_y^2(\tau)$ does not converge for white PM noise and flicker PM noise, unless a lowpass filter is introduced to limit the bandwidth, whose cutoff frequency is denoted with f_H .

4.0.3 Overlapped Allan Variance

An efficient way to measure the Allan variance is to sample $x(t)$ at the rate $1/\tau_0$, taking $\tau = m\tau_0$, integer m . The k -th value of the fractional frequency is evaluated as

$$\bar{y}_k = \frac{x_{k+m} - x_k}{m\tau_0}. \quad (4.7)$$

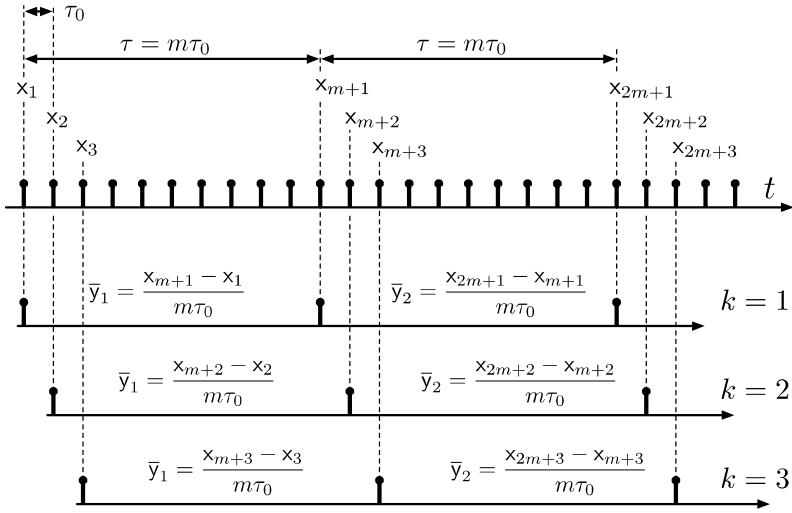


Figure 4.1: Evaluation of the overlapped Allan variance.

The overlapped Allan variance consists of using partially overlapped realizations of $\bar{y}_2 - \bar{y}_1$ in (4.2), separated by the minimum amount τ_0 . This concept is illustrated in Fig. 4.1. Accordingly, (4.4) becomes

$$\tau = m\tau_0$$

$$\widehat{\sigma_y^2}(m\tau_0) = \frac{1}{2(M-1)} \sum_{k=1}^{M-1} \left[\frac{x_{2m+k} - 2x_{m+k} + x_k}{m\tau_0} \right]^2, \quad (4.8)$$

which takes $2m + M - 1$ samples, thus a measurement time $\mathcal{T} = (2m + M - 2)\tau_0$.

A first advantage of overlapping is smaller uncertainty. In fact, the confidence interval $\Delta\sigma_y/\sigma_y = \sqrt{2/\mathfrak{D}}$ is related to the number \mathfrak{D} of degrees of freedom. In turn, \mathfrak{D} is equal to the number of samples of $\bar{y}_{k+1} - \bar{y}_k$ in the case of white PM noise (uncorrelated samples), it gets progressively smaller for slower noise phenomena, and it degenerates to 2 in the case of pure drift. A simple example deserves attention. Suppose we have a record of 3×10^5 samples spaced by $\tau_0 = 100$ ms, thus $\mathcal{T} = 3 \times 10^4$ s = 8.33 hours acquisition time. At $\tau = 10^4$ s, we have only 2 realizations of $\bar{y}_{k+1} - \bar{y}_k$ using (4.4), while in the same conditions we have 10^5 realizations if we opt for (4.8). At a deeper sight such realizations are highly correlated, but the larger number is still beneficial to the reduction of uncertainty.

A further advantage of overlapping is that it solves the erratic response of AVAR in the presence of cyclic disturbances of period $T \approx \tau/0.45$ (2.2τ), like the diurnal temperature. This happens because multiple realizations of AVAR are averaged, based on measures of $\bar{y}_2 - \bar{y}_1$ progressively shifted by $m\tau_0$, as shown in Fig. 4.1.

The Allan variance with no overlap is seldom used, if ever. Generally, the term ‘Allan variance’ refers to the overlapping algorithm with no need of saying so.

4.0.4 Frequency Counters and Weighted Averages (Region 2.1)

Introducing the classical Allan variance, we have defined \bar{y} as the uniform average of $y(t)$ over the time τ . Other options make sense, based on the redefinition of the average \bar{y} as

$$\bar{y}(\tau) = \int_0^\infty y(t) w(t; \tau) dt. \quad (4.9)$$

The weight function $w(t; \tau)$ takes different forms, among which the following deserve attention

- $w_\Pi(t; \tau)$ gives the uniform average, the same used in Sec. 16.4 (there is still no change). The classical reciprocal counter, also called Π counter after [Rubiola, 2005b], uses w_Π . The symbol Π recalls the rectangular shape of $w_\Pi(t; \tau)$.
- $w_\Lambda(t; \tau)$ gives the triangular average, which is calculated averaging on a sequence of highly overlapped rectangular averages. The corresponding instrument is the Λ counter after [Rubiola, 2005b]. Of course, the Greek letter Λ is chosen because of its triangular shape. The benefit of such counter is a high rejection of the wideband white PM noise of the trigger at the counter input. Some commercial instruments implement the Λ averaging, often without saying. They can be identified from the ‘precision’ (response to the trigger noise) proportional to $1/\tau\sqrt{\tau}$ instead of $1/\tau$.
- $w_\Omega(t; \tau)$ relates to a frequency measurement implemented as a linear regression on phase-time data. The corresponding instrument is called Ω counter after [Rubiola et al., 2016]. The benefit of the Ω counter is the highest rejection of the white PM noise, by theorem. The Greek letter Ω is the graphically closest to the parabolic shape of the frequency response. Besides, the letter Ω indicates that this is the ultimate counter, to the extent that no other counter performs better rejection of white PM noise.

The ‘precision’ (response to the trigger noise) is proportional to $1/\tau\sqrt{\tau}$ as in the Λ counter, just with 1.25 dB lower background noise (a factor of 3/4). Very few commercial instruments implement this algorithm.

- $w_\Delta(t; \tau)$ is equivalent to the difference between two contiguous measures taken with a Π frequency counter. This option is listed only for completeness, because to the best of our knowledge it is not implemented in commercial counters.

Like the overlapped Allan variance, $y(t)$ is sampled at a suitable frequency $1/\tau_0$, thus it holds that $\tau = m\tau_0$. As a consequence, the patterns of the Λ and Ω counters shown in Region 2.1 hold for $m \gg 1$, or equivalently $\tau \gg \tau_0$, so that the continuous approximation holds.

4.0.5 Generalized Two-Sample Variances (Region 2.2)

The definition (4.2) is more general than the classical Allan variance. In fact, feeding the weighted averages (4.9) into (4.2) results in different types of variance, which can

all be described by the same formula

$$\sigma_y^2(\tau) = \mathbb{E} \left\{ \int_0^\infty \left[y(t) w(t; \tau) \right]^2 dt \right\}. \quad (4.10)$$

Notice that the difference $\bar{y}_2 - \bar{y}_1$ is now included in the wavelet-like function $w(t; \tau)$. The latter is similar to a wavelet, but for the normalization for finite-power signals (power-type signals) instead of finite-energy signals (energy-type signals).

We have the following options

$$\begin{aligned} w_A(t; \tau) &= w_\Pi(t - \tau; \tau) - w_\Pi(t; \tau) && \rightarrow && \text{AVAR } {}^A\sigma_y^2(\tau) \\ w_M(t; \tau) &= w_\Lambda(t - \tau; \tau) - w_\Lambda(t; \tau) && \rightarrow && \text{MVAR } {}^M\sigma_y^2(\tau) \\ w_P(t; \tau) &= w_\Omega(t - \tau; \tau) - w_\Omega(t; \tau) && \rightarrow && \text{PVAR } {}^P\sigma_y^2(\tau) \\ w_H(t; \tau) &= w_\Delta(t - \tau; \tau) - w_\Delta(t; \tau) && \rightarrow && \text{HVAR } {}^H\sigma_y^2(\tau). \end{aligned}$$

As with the frequency counters, the patterns of $w_M(t; \tau)$ and $w_P(t; \tau)$ shown in Region 2.2 are the limit for $\tau \gg \tau_0$, or equivalently for $m \gg 1$, where the continuous approximations holds. The formulae of Region 2.7 hold under this assumption.

Using commercial frequency counters to get a data stream \bar{y}_k , we recommend attention to the contiguity of averages. Such contiguity is implied in the waveforms of Region 2.2. The case of MVAR is subtle because ‘contiguous’ triangular averages overlap by one side, instead of touching one another by one edge. Something like $\text{///}\dots$ instead of $\text{^}\text{^}\text{^}\text{^}\dots$. So, if the measurement of \bar{y} takes 2τ , the measurement of $\bar{y}_2 - \bar{y}_1$ takes 3τ . Reference [Dawkins et al., 2007] points out that some Keysight counters provide a data stream of the latter type, where continuous averages touch one another by one edge.

Our notation ${}^A\sigma_y^2(\tau)$, ${}^M\sigma_y^2(\tau)$, ${}^P\sigma_y^2(\tau)$ and ${}^H\sigma_y^2(\tau)$, seems more elegant than AVAR, MVAR, PVAR and HVAR, but it is not used in the literature.

We draw the reader’s attention to the fact that (4.10) may be misleading in the case of HVAR because it hides the fact that HVAR is a second-difference variance. HVAR is rather different from the other variances, chiefly in the fact that it converges for integrated flicker FM and random run FM ($1/f^3$ and $1/f^4$ FM), and it is blind to frequency drift.

4.0.6 The Time Variance TVAR (Region 2.6)

The Time Variance TVAR is defined as

$$\sigma_x^2(\tau) = \frac{1}{3} \tau^2 {}^M\sigma_y(\tau). \quad (4.11)$$

TDEV, the square root of TVAR, is an estimator of the uncertainty of the time elapsed after the duration τ . TDEV is often used in telecom, e.g. to assess the Time Interval Error TIE, defined in [ITU Working Group 15, 2000], and the holdover performance, i.e., the clock error accumulated between synchronizations [ETS, 2002, Part 5.1: Timing characteristics of slave clocks...].

4.0.7 Additional Options

The following options are highly specialized instances of the Allan variance, beyond the scope of this article. They are mentioned for completeness, and left to experts.

The Dynamic Allan Variance

A time series of N data is sliced into m sub-series of $n = N/m$ data. Computing the AVAR for each sub-series, we end up with a 3D plot which shows the changes of AVAR vs time, most useful for diagnostic purposes. This variance and its properties is found in a series of articles by the same team [Galleani and Tavella, 2009, Galleani, 2010, Galleani, 2011, Galleani and Tavella, 2015, Galleani and Tavella, 2016], but it does not seem to have been followed by other authors and research teams.

The Total Variance

The time series is circularized by joining a copy with time reversed, as often done in the domain of spectral analysis. Circularization cannot increase the number \mathfrak{D} of degrees of freedom inherent in the experimental outcomes, but it makes their exploitation more efficient for the detection of certain phenomena. This concept, described in [Greenhall et al., 1999, Howe, 2000], can be applied to all wavelet variances.

Thêo1, ThêoH and ThêoBR

These estimators are based on the idea that two measures of duration $\tau' \ll \tau$ whose centers are spaced by τ provide a precise estimation of the slow processes occurring at τ , under the condition that τ' is long enough to average out the fast processes. These ideas are found in [Howe and Pepler, 2003, Howe, 2006, Taylor and Howe, 2010]. The benefit is to extend the maximum τ beyond $\mathcal{T}/2$, where \mathcal{T} is the duration of the data record. For example, applying ThêoH to a long data record of an atomic time scale, it is possible to extend the plot up to $\tau = 0.8\mathcal{T}$. So, these variances are even more efficient than AVAR at estimating the fluctuations occurring at large τ . This benefit comes at the cost of higher computing burden and additional interpretation difficulty.

4.0.8 Choosing the Most Appropriate Variance

For general use, the variances described are broadly similar to one another, and none is really “the best” or just “bad.” Each one has its own ‘personality,’ which makes it more suitable in some specific case, and weaker in other cases. Such personality follows from the wavelet-like patterns shown in Region 2.2. A summary of the main options is given below.

Normalization

In signal processing, normalization for white noise is the most common option. A different choice is made here because the two-sample variances are issued from time-keeping. The normalization is chosen for all the variances to have the same response $\sigma_y^2(\tau) = \frac{1}{2}D_y^2\tau^2$ to the linear frequency drift D_y (the “linear drift” row in the lower part of Region 2.7). A consequence is that different variances have different responses to the other noise processes because the spectral response $|H(f\tau)|^2$ differs (Region 1.5).

Additionally, the corners separating the noise processes are not the same (Regions 1.7, and 2.3–2.5). Comparing the plots requires a small effort of interpretation. For example, the white FM noise $S_\varphi = \mathbf{b}_{-2}/f^2 \rightarrow S_y = \mathbf{h}_0$ shows up as

$$^A\sigma_y^2(\tau) = \frac{1}{2} \frac{\mathbf{h}_0}{\tau} = 0.50 \frac{\mathbf{h}_0}{\tau} \quad (\text{Allan})$$

$$^M\sigma_y^2(\tau) = \frac{1}{4} \frac{\mathbf{h}_0}{\tau} = 0.25 \frac{\mathbf{h}_0}{\tau} \quad (\text{Modified Allan})$$

$$^H\sigma_y^2(\tau) = \frac{1}{3} \frac{\mathbf{h}_0}{\tau} \simeq 0.33 \frac{\mathbf{h}_0}{\tau} \quad (\text{Hadamard})$$

$$^P\sigma_y^2(\tau) = \frac{3}{5} \frac{\mathbf{h}_0}{\tau} = 0.60 \frac{\mathbf{h}_0}{\tau} \quad (\text{Parabolic}),$$

with a maximum difference of a factor of 2.4 (3.8 dB).

AVAR

This is the best choice when we want to evaluate $\sigma_y^2(\tau)$ up to the largest τ for a given data record of duration \mathcal{T} , that is, $\tau = \mathcal{T}/2$. This is the typical case of atomic time scales, where the oscillators are continuously monitored, and we focus on the slow processes.

Because $w_A(t; \tau)$ takes 2τ for one realization of $\sigma_y^2(\tau)$, averaging on $M \gg 1$ realizations is made possible by overlapping the measures with $\tau_0 \ll \tau$ (Sec. 16.4).

The uniform weight of $w_A(t; \tau)$ features the highest efficiency in picking up the energy of y . By contrast, AVAR is unsuitable to the measurement of white PM noise because $^A\sigma_y^2(\tau) \simeq 0.076 f_H \mathbf{h}_2 / \tau^2$ (Region 2.7), thus the result is highly dependent on the bandwidth f_H . This is not a problem for the slow phenomena we mentioned. Finally, it is worth mentioning that, in the presence of white FM noise only, AVAR gives the same result as the classical variance.

MVAR

This is a good choice in the presence of wideband noise typical of fast processes. MVAR originates from optics, where a precise and efficient measurement of white PM noise is a desired feature.

By contrast, MVAR is inferior to AVAR in the efficient use of \mathcal{T} because the support of $w_M(t; \tau)$ is 3τ wide instead of 2τ . This may not be a problem when the physical phenomena we are interested in occur at small or moderate τ , say hours.

HVAR

This variance is useful for the measurement of strong slow phenomena occurring in some circumstances, for example in the absence of temperature stabilization. In fact, unlike the other variances described here, it converges for integrated flicker FM noise and for integrated random walk FM noise, also called ‘random run.’ Such processes are the \mathbf{h}_{-3}/f^3 and \mathbf{h}_{-4}/f^4 terms of $S_y(f)$, or equivalently the \mathbf{b}_{-5}/f^5 and \mathbf{h}_{-6}/f^6 terms of $S_\varphi(f)$. By contrast, HVAR is blind to linear drift. This makes HVAR a specialized tool, particularly useful when high drift makes it difficult to estimate the other noise parameters.

Finally HVAR, like AVAR, gives ambiguous response to white PM noise, as it depends on the instrument bandwidth f_H . Like MVAR, $w_H(t; \tau)$ takes a time equals to 3τ .

PVAR

The computation of $\sigma_y^2(\tau)$ at $\tau = m\tau_0$ requires a data record of duration $\mathcal{T} = k\tau$, where $k > 2$ for AVAR and PVAR, and $k > 3$ for MVAR and HVAR. The minimum \mathcal{T} depends on the noise process, on the confidence level required, and on the variance we choose. Running the measurement, we start seeing white PM noise at short \mathcal{T} , then flicker PM, white FM etc. as \mathcal{T} increases. Now we take a different standpoint, asking which is the minimum \mathcal{T} to detect a ‘new’ noise phenomenon, out of the ‘previous,’ faster one. For example, which is the shortest \mathcal{T} to see that flicker FM is above the white FM, with 95 % probability? Among the four variances considered here, PVAR is the best at doing this from white PM to RW FM. Yet, the corner between RW FM and frequency drift is still better detected by AVAR.

4.0.9 Some Pieces of Advice

Beginners should restrict their attention to ADEV and MDEV, the square root of AVAR and MVAR. ADEV and MDEV are both available in commercial instruments, and both benefit from the size effect of a wide community.

For historical reasons, ADEV is definitely the manufacturers’ preferred option. Reading technical documentation, we recommend attention to a possible confusion between ADEV and MDEV under the term ‘Allan deviation,’ with a possible ‘modified’ omitted or implied. The ITU-T Recommendation G.8260 [ITU Working Group 15, 2000] compares ADEV, MDEV and TDEV from the standpoint of telecommunications. MDEV seems the favorite tool in telecommunications [Bregni, 2016], and enables the direct calculation of TDEV using (4.11). However, MDEV takes 50% longer acquisition time \mathcal{T} .

Doing one’s own measurements, MDEV is in most cases the best compromise. In fact, MDEV improves on ADEV in the detection of fast noise processes (white and flicker PM), and is as suitable as ADEV to detect all the other noise processes. AVAR is still the best option for the measurement of the atomic clocks² intended for time scales, where increasing \mathcal{T} is costly or impossible. Finally, the reader should remember that normalization makes MVAR always ‘optimistic’ compared to AVAR, as seen in the related columns in Region 2.7.

Looking at the future, PDEV may replace MDEV because it outperforms it in all parameters at no cost but computing power. Likewise, Théo may replace ADEV because of the more efficient use of \mathcal{T} at representing longer values of τ .

²In the international coordination of metrology under the guidance of BIPM, an atomic oscillator is a ‘clock’ only after 6 months of uninterrupted contribution to UTC/TAI. Some commercial oscillators are called clock by the manufacturer implying that they are suitable to contribute to UTC/TAI. This is perfectly sound because hundreds of such Cesium clocks actually do this — and Hydrogen masers as well. However, the term clock is also used as a pretentious replacement for a precision oscillator even if reliability and long-term stability are insufficient for the oscillator to be even considered as a contributor to UTC/TAI.




Figure 4.2: Example of MDEV obtained with the SigmaTheta software tool. Notice the difference between the calculated values of MDEV (black crosses) and the Bayesian estimates (green donuts), and the asymmetry of the uncertainty bars. Data are courtesy of Moustafa “Mouss” Abdel Hafiz, FEMTO-ST Institute.

4.0.10 Example of MDEV

We measure the stability of a miniature Cs oscillator. The oscillator is a laboratory prototype based on the Coherent Population Trapping (CPT) principle in a Cs microcell. The actual experiment is described in [Abdel Hafiz et al., 2022], and the fabrication of the microcell in [Hasegawa et al., 2011]. For technical reasons specific to the experiment, the frequency is sampled at 150 ms intervals and pre-processed to provide a stream of values of $y_k s$ uniformly averaged on contiguous intervals of 1 s. This sets $\tau_0 = 1$ s. The total number of samples is 389998, for a total duration $\mathcal{T} \approx 4.5$ days.

Fig. 4.2 shows the MDEV of the experiment described, processed with the SigmaTheta software tool. Starting from uniformly averaged data, a practical minimum of 8–10 samples is needed to approximate the triangular average (Regions 2.1 and 2.2.). Thus, MDEV is plotted for $\tau = 2^n \tau_0$ starting from $\tau = 8\tau_0 = 8$ s. At the scale of this experiment, the H maser used as the reference can be considered ideally stable, and the noise of the instrument is negligible as well.

The SigmaTheta software package provides the following pieces of information (see Sections 4.2 and 4.3)

- The bare values of $M_{\sigma_y}(\tau)$, shown as black crosses,
- The Bayesian estimates of $M_{\sigma_y}(\tau)$, shown as green donuts,

- The uncertainty bars,
- The identification of the most relevant noise processes, shown as the colored straight lines of slope $1/\sqrt{\tau}$, constant vs τ , and $\sqrt{\tau}$.

It is worth pointing out that the uncertainty bars systematically extend upwards more than downwards. This is a kind of ‘signature’ of the inverse problem, as opposed to the bare simulation approach.

4.0.11 Suggested Readings About Variances

General references

The free booklet [Riley,] is probably where most readers should start. Sponsored and distributed by NIST, it provides an extensive coverage of most variances (AVAR, MVAR, HVAR etc.) and the evaluation of the confidence intervals, with numerous examples and plots made with Stable32. Reference [Stein, 2010] is a review article about the Allan variance, and [Enzer et al., 2021] suggests that the Allan variance can be used as a diagnostic tool. A wealth of information is available in a Special Issue of the IEEE Transact. UFFC celebrating the 50th anniversary of the Allan Variance [Levine et al., 2016]. An historical review is available [Allan and Levine, 2016], written by two of the most important contributors to the rise of this branch of knowledge. Reference [Allan, 1966] is the original article that introduces the sample variances, later called Allan variance, and [Allan and Barnes, 1981] introduces MVAR.

Π and Λ counters, and the related statistics

Reference [Rubiola, 2005b] is the first article which defines Π and Λ counters and the mathematical framework underneath, later extended to the case of non-overlapping triangular averages [Dawkins et al., 2007]. However, the basic ideas were already in [Kramer and Klische, 2001, Kramer and Klische, 2004], yet without developing the statistical framework. A wealth of practical knowledge about the architecture of high-resolution counters is available in a review article [Kalistz, 2004].

The Ω counter and the Parabolic Variance

The linear regression of phase data is a rather obvious way to estimate a frequency. It was used in the HP5371A HP5372A time interval analyzers³ in the late 1980s, and at Pendulum [Johansson, 2005]. The name “ Ω counter” comes from [Rubiola et al., 2016], which introduces the related mathematical framework for frequency metrology. Feeding the linear-regression estimates into (4.2), we get PVAR. This idea came independently from [Benkler et al., 2015] and [Vernotte et al., 2016], and the name PVAR was decided together by the two teams. Reference [Vernotte et al., 2016] digresses the advanced statistical properties of PVAR, including the Bayesian statistics and the minimum duration of the data record to detect a noise process.

³Information provided by Magnus Danielsson, NetInsight, Sweden.

Aliasing

Reference [Vernotte et al., 1998b] provides theory and insight on spectral aliasing, and [Calosso et al., 2016] gives an interesting perspective about aliasing in the AVAR and MVAR, covering the effect of spectral bumps and ‘blue noise,’ often found in long-range optical frequency-distribution systems. Finally, [Bernier, 1987] provides useful insight in aliasing and cutoff frequency for MVAR.

4.1 Conversions

4.1.1 Visual Inspection of Plots (Regions 1.6–1.7)

Regions 1.6–1.7 show a plot of $S_y(f)$ beside AVAR $\sigma_y^2(\tau)$ for the noise processes from white PM to random-walk FM. The following facts deserve attention.

- There is a kind of mirror symmetry between the plots of $S_y(f)$ and $\sigma_y^2(\tau)$. The fastest process, white PM, is on the right-hand side of $S_y(f)$, and on the left-hand side of $\sigma_y^2(\tau)$. Vice versa, the FM random walk is on the left-hand side of $S_y(f)$, and on the right-hand side of $\sigma_y^2(\tau)$.
- The cutoff frequency f_H has a dramatic effect on white PM noise, only a weak effect on the flicker PM noise, and virtually no effect on slower processes.
- The $1/\tau^2$ region of $\sigma_y^2(\tau)$ is ambiguous, in that it represents both white PM and flicker PM noise.
- The conversion from $S_\varphi(f)$ to $\sigma_y^2(\tau)$ is always possible, while the opposite suffers from limitations. This is emphasized by the road signs between Regions 1.6 and 1.7.
- The corner τ where the processes cross one another may occur rather far from the values one expects intuitively.

If we have both phase noise and variance measures, the conversion from $S_\varphi(f)$ to $\sigma_y^2(\tau)$ is a great way to check on consistency. However, this is possible only if $S_\varphi(f)$ extends to sufficiently low f to reveal the slow processes shown by $\sigma_y^2(\tau)$. Regardless, the conversion is an exercise of interpretation we recommend.

4.1.2 Conversion from PM Noise to Allan Variance (Region 2.7)

The reader should first refer to Region 1.11 for the conversion $S_\varphi(f) \rightarrow S_y(f)$. Then, to Region 2.7 for $S_y(f) \rightarrow \sigma_y^2(\tau)$, based on (4.5)-(4.6). A simple procedure is shown in Fig. 4.3, and detailed below.

0. Start with a log-log plot of phase noise. Since in most cases the phase noise is given as $10 \log_{10}[\mathcal{L}(f)]$, you have to convert it into $S_\varphi(f)$ using $S_\varphi(f) = 2\mathcal{L}(f)$, i.e., add 3 dB (Region 1.11)
1. Approximate the true spectrum with the straight lines which represent the polynomial law. Proceed from the right-hand side (white PM noise) to the left, not

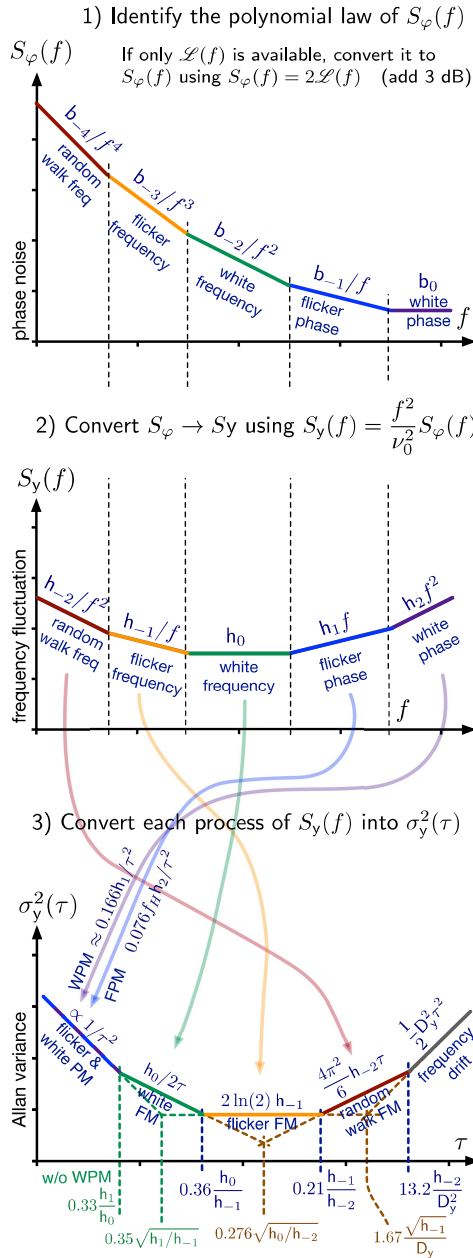


Figure 4.3: Conversion from phase noise to Allan variance. The colors of the straight line approximation recall the frequency, from reddish (low) to bluish (high). The conversion from phase noise to the other two-sample variances is an obvious extension.

vice versa. This best done with a drawing app inserting a straight line of exact slope (0, -1 , -2 etc.) and shifting it to fit the plot. The coefficients b_0 , b_{-1} , b_{-2} , etc. are the value of the corresponding straight line found at $f = 1$ Hz. Even with little training, visual inspection is efficient at approximating the noise process and ignoring the artifacts. Sliding old-fashioned set drafting squares on a printed spectrum is a good alternative to a drawing app. Worked-out examples are available in [Rubiola, 2010, Chapter 6].

2. Convert $S_\varphi(f)$ into $S_y(f)$ using

$$h_n = \frac{1}{\nu_0^2} b_{n-2},$$

which is equivalent to $S_y(f) = \frac{f^2}{\nu_0^2} S_\varphi(f)$ (Region 1.11). Newcomers may be surprised by the small value of the h_n coefficients, due to ν_0^2 in the denominator.

3. Sketch the Allan variance using the pattern of Region 1.7 (right) and the formulas found in the ‘AVAR’ column of the Table in Region 2.7. Each process (white PM, flicker PM, etc.) requires its own formula, found in Region 2.7.

It goes without saying that the method described also applies to the other two-sample variances discussed in Sec. 16.5, just picking up the appropriate column in Region 2.7. The extension to non-integer slopes is found in [Walter, 1994, Vernotte et al., 2021].

4.1.3 Variance to Spectrum Conversion (Regions 1.6, 1.7, 2.7)

The variance to spectrum conversion is not possible in the general case, but we can get useful information using the formulas given in Region 2.7 assuming that the spectrum is smooth and follows the polynomial law (see [Greenhall, 1998]).


A first limitation is that in the case of AVAR and HVAR, both white and flicker PM show up as $\sigma_y^2 \propto 1/\tau^2$, thus they cannot be divided. Additionally, AVAR and HVAR are almost unusable in the white PM region because $\sigma_y^2 \propto f_H$. The cutoff f_H does not appear explicitly in MVAR, PVAR and TVAR. However, the effect of f_H is hidden in the sampling process itself and then in the data. Since data are sampled at τ_0 interval, the Nyquist frequency is $\frac{1}{2\tau_0}$. Thus, if $f_H > \frac{1}{2\tau_0}$, there will be spectral aliasing and the white noise level will be overestimated.

A second limitation is related to resolution. In fact, the inherent resolution of $\sigma_y^2(\tau)$ is one octave in τ , as it follows from the bandwidth of the main lobe in Region 1.5. Thus, $\sigma_y(\tau)$ is usually plotted for τ in geometric series like 1, 2, 4, 8... Conversely, the frequency resolution of $S_\varphi(f)$ depends on the acquisition time, and plots usually represent $S_\varphi(f)$ with a resolution of 50–100 points/decade, that is, 15–30 points per octave.

4.1.4 The Cutoff Frequency f_H and the Sampling Interval τ_0

The cutoff frequency f_H , extensively used in the spectrum-to-variance conversions is often a source of confusion. Notice that an anti-aliasing filter is necessary, otherwise

(a)



(b)




Figure 4.4: Beat method, applied (a) to electrical signals, and (b) to optical signals. See Section 16.8.5 for the general description.

the variance does not converge in the presence of white and flicker PM noise [Bernier, 1987].

The phase time $x(t)$ has finite bandwidth which results from the measurement process and from the architecture of the instrument. This is generally described as the noise equivalent bandwidth, denoted with f_H .

The variance is evaluated after sampling $x(t)$ at an appropriate frequency $1/\tau_0$, which requires that $f_H < 1/(2\tau_0)$. If this condition is not met, aliasing takes place. The white PM noise is folded $2f_H\tau_0$ times to the first Nyquist zone, and the observed white noise level is $2f_H\tau_0k_0$ instead of k_0 . Flicker noise too is subject to aliasing, even though the impact on the results may be smaller. References [Bernier, 1987, Vernotte et al., 1998a, Vernotte et al., 1998b] detail the specific problems related to our domain.

4.1.5 Example (f_H)

Let us take an example from optics, where we beat two 1550-nm lasers in a photodiode with the scheme of Fig. 4.4, getting a RF tone at $\nu_b = 60$ MHz. For technical reasons we choose to filter such signal with ± 0.5 MHz bandpass centered at ν_b . The filter halfwidth is 2.6×10^{-9} of the optical carrier, thus the laser instability must be $< 5 \times 10^{-10}$ for ν_b to be decently centered in the filter pass band. The noise bandwidth is equal to the filter half-width, i.e., $f_H = 500$ kHz.

First option

we measure the beat note with a counter sampling at $\tau_0 = 1$ ms interval (for example, the old good K&K counters, or the more recent version made by Lange Electronic). In the conditions described, aliasing increases the white PM noise level by a factor $2f_H\tau_0 = 1000$.

Second option

we measure the same beat note with an instrument based on direct digitization (Sec. 16.8.8 and Fig. 16.9) or a Tracking DDS (Sec. 16.8.9 and Fig. 16.10). Inside the instrument, the ADCs are preceded by antialiasing filters, and a lowpass reduces the bandwidth of $x(t)$ after detection. Thus the condition $f_H < 1/(2\tau_0)$ is met, and there is no aliasing. In some instruments f_H can be set by the user.

4.2 Confidence Intervals

Assessing the confidence intervals is all about understanding the interplay between the ‘true value’ $\sigma_y^2(\tau)$ and its estimate $\widehat{\sigma_y^2}(\tau)$. The latter is evaluated with (4.3), or equivalently (4.4) for AVAR, and with similar formulae for other variances. Since everything in this Section applies separately to each value of τ , we let τ implied in $\widehat{\sigma_y^2}$ and σ_y^2 .

4.2.1 Direct Problem

The *direct problem* is similar to computer simulations, where we add noise to a deterministic phenomenon. Without noise, the right-hand side of (4.3) gives σ_y^2 . In this case, the term ‘true value’ for σ_y^2 is fully legitimate because uncertainty is zero. Introducing the noise, the same equation gives the estimate $\widehat{\sigma_y^2}$, together with the associated Probability Density Function, or PDF, $p(\widehat{\sigma_y^2}|\sigma_y^2)$. The notation $p(|)$ stands for conditional probability, and the vertical bar ‘|’ reads ‘given’ or ‘knowing,’ as in $p(\widehat{\sigma_y^2})$ given $p(\sigma_y^2)$.

4.2.2 Inverse Problem

The simulation approach does not help to grab the hidden reality σ_y^2 from the outcome $\widehat{\sigma_y^2}$. The answer is the *inverse problem*, which consists of inferring the true σ_y^2 and the associated PDF $p(\sigma_y^2|\widehat{\sigma_y^2})$.

Otherwise stated, the direct problem targets the ‘cause \rightarrow effect?’ relationship, where ‘?’ emphasizes the unknown. In contrast, the inverse problem targets ‘cause? \leftarrow effect’. Interestingly, the ‘true value’ σ_y^2 in the direct problem and the ‘inferred true value’ σ_y^2 in the inverse problem are different things, but they are generally denoted with the same symbol in the literature.

The evaluation of the confidence intervals in the inverse problem may be too complex for some readers. These readers may keep only the two results summarized below,

skip Section 4.2.3, and go straight to the choice of the appropriate software package.

- When a small number $M - 1$ of realizations $\bar{y}_{k+1} - \bar{y}_k$ is available, the ‘error bars’ are large and asymmetric. The half bar directed upwards is wider than the downwards half bar. This is seen in the example Fig. 4.2 for $\tau \geq 4000$ s, and qualitatively illustrated in Region 1.7.
- With large M , there is no practical difference between direct problem and inverse problem. The confidence intervals are small and symmetric, approaching $1/\sqrt{M-1}$, as seen in Fig. 4.2 for $\tau \leq 2000$ s.

4.2.3 Application of the Bayesian Statistics to Allan Variances

Everything starts with the Bayes theorem, which states that

$$p(\theta|\xi) = \frac{\pi(\theta)p(\xi|\theta)}{\pi(\xi)}, \quad \pi(\xi \neq 0), \quad (4.12)$$

where $p(\cdot)$ denotes the *posterior* PDF, and $\pi(\cdot)$ denotes the *prior* PDF⁴. It is worth mentioning that some authors use the same symbol for both prior probability and posterior probability. In the measurement of frequency stability, the experimental value $\widehat{\sigma_y^2}$ is identified with ξ , and the unknown ‘true’ σ_y^2 with θ . Thus, we have to infer a confidence interval on σ_y^2 .

The central limit theorem suggests that the \bar{y}_k are Gaussian, as they result from a lot of data. Thus, we assume that their differences are Gaussian centered. Thus, (4.3) indicates that $\widehat{\sigma_y^2}$ is described by a χ^2 distribution with \mathfrak{D} degrees of freedom⁵. Such distribution is denoted with $\chi_{\mathfrak{D}}^2$. Of course, with M values of \bar{y}_k it holds that $\mathfrak{D} \leq M - 1$, where the equality indicates that all the terms $(\bar{y}_{k+1} - \bar{y}_k)$ of the sum are statistically independent. Greenhall and Riley provides a very useful method to evaluate \mathfrak{D} [Greenhall and Riley, 2003].

Since the random variable ξ is $\chi_{\mathfrak{D}}^2$ distributed, the cumulative density function (CDF) of ξ knowing θ , denoted with $F(\xi|\theta)$, is also known in analytic form. The inverse CDF, available in the major mathematical libraries, enables the user to compute the confidence interval.

The above is for the direct problem. The inverse problem can be solved thanks to the relevant property that a $\chi_{\mathfrak{D}}^2$ distribution is defined by one and only parameter, \mathfrak{D} . It has been proved that such distributions are “fiducial” distributions [Fisher, 1935], which means that the equality

$$F(\theta|\xi) = 1 - F(\xi|\theta) \quad (\text{fiducial}) \quad (4.13)$$

holds in both frequentist inference and Bayesian inference, provided that a $1/\theta$ prior (prior of total ignorance) is chosen [Lindley, 1958, Vernotte and Lantz, 2012]. This

⁴More specifically, $\pi(\theta)$ is the a-priori knowledge before any measurement, and $\pi(\xi)$ is an unknown function we don’t care about because it is independent of θ . The prior $\pi(\xi)$ can be determined using the property that $\int p(\theta|\xi) d\theta = 1$.

⁵In the literature about statistics the degrees of freedom are more often denoted with ν , but in our notation ν is used for the carrier frequency.

implies that the confidence intervals given by frequentist or Bayesian methods are the same and are easy to compute.

Note that, because the $\chi^2_{\mathfrak{D}}(x)$ distribution is strongly asymmetric with steep rise at small x and slow decay at high x , it results from (4.13) that $p(\theta|\xi)$ must have a steep side at some high x and a slow decay towards $x = 0$. The consequence is that the error bars on a log-log plot of Allan variance are extend upwards more than downwards. This behavior is more remarkable at small \mathfrak{D} .

4.3 Software Tools

A quick Internet search reveals that there is a rather broad choice of software packages for the Allan variance, but almost none for phase noise. This relates to the fact that the variances are used to investigate slow phenomena, thus they require rather low sampling rate and small digital storage space. A value of 1 kS/s is generally sufficient for all practical cases, but 1 S/s or 10 S/s is most often used in atomic time scales. Conversely, it is quite common to plot the phase noise up to 1 MHz Fourier frequency. This requires a practical minimum of 2.5 MS/s, allowing a mere 250 kHz for the anti aliasing filter to roll off. In turn, a transfer rate of 10 MB/s is necessary for a cross-spectrum system under the hypothesis that the phase is encoded on 16 bits. Thus, a measurement lasting 100 s takes 1 GB disk space. To the best of our knowledge, noise analyzers do not save or transfer raw data of this size and at this rate. TimeLab⁶ is no exception, to the extent that continuous gap-free IQ data are transferred from the 3120A/5330A/53100A instruments to the computer at reduced bandwidth (100 kHz). Faster IQ data required to plot spectra up to the maximum frequency of 1 MHz are transferred in bursts with 10% duty cycle. Only the spectra bins are saved in the .TIM files, together gap-free data at significantly lower rate (1 kS/s) for time measurements and Allan variances.

We present a selection of software packages (Table 4.1), chosen for their scientific value or for their wide use. The grayed area highlight some interesting features.

4.3.1 Features

Main purposes

AllanTools, SigmaTheta and Stable32 are intended for data analysis. By contrast, TimeLab is a tool for data acquisition, with limited analysis capabilities. It supports the phase noise analyzers from Jackson Lab and Microchip (formerly Microsemi), and a few frequency counters.

Graphical interface versus scripting

Albeit elderly, Stable32 has an efficient graphical interface, which makes it a great choice for occasional users. For this reason, it is by far the most widely used. AllanTools and SigmaTheta require programming skills. On the other hand, scripting is great in that it enables the analysis of a bulk of data sets at once.

⁶Information provided by John Miles, September 2022.

Table 4.1:
SOFTWARE TOOLS FOR THE TWO-SAMPLE VARIANCES

		Tool			
		AllanTools	SigmaTheta	Stable32	TimeLab
Features ¹					
Fractional frequency	ADEV	Y ²	Y	Y ²	Y
	GCODEV	beta	Y	–	–
	GRADEV	Y	–	–	–
	HDEV	Y ²	Y	Y ²	Y
	MDEV	Y	Y	Y	Y
	PDEV	–	Y	–	–
	Thêo1	Y	–	Y	–
	ThêoH	–	–	Y	–
	TOTDEV	Y	–	Y	–
Time	MTIE	Y	–	Y	Y
	TDEV	Y	–	Y	Y
	TIErms	Y	–	Y	–
Other	Estimator	direct	Bayes	direct	direct
	Error bars	LA & GR ^{3,4,5}	GR ⁴	GR ^{4,5}	LA ³
	Curve fit	–	Y	–	–
General	Type	Python lib	CL ⁶ , API ⁶ , GUI ⁶	GUI ⁶	GUI ⁶
	OS ⁷	L/M/W	L/M ⁸ /W ⁸	L ⁹ /W	W
	License	LGPLv3+ (open)	CeCILL (open)	© ¹⁰ (free)	© (free)
(1) Here, “DEV” implies that both DEV and VAR are available					
(2) Also, non-overlapped algorithm					
(3) LA = Lesage-Audoin algorithm, based on stdev(σ^2)					
(4) GR = Greenhall-Riley algorithm, based on $\chi^2(\sigma^2)$					
(5) See main text					
(6) CL = Command line, API = Application Programming Interface, GUI = Graphical user interface					
(7) L = Linux, M = macOS, W = Windows					
(8) API and GUI on Mac and Windows require recompiling					
(9) Wine environment, or Windows virtual machine					
(10) The team is working at cleaning the code for public release					

Mathematical functions

All the packages provide ADEV, MDEV and the other mostly-used functions. For simplicity, DEV in Tab. 4.1 stands for DEV or VAR. When it comes to more exotic functions like Th  o and the Gros Lambert covariance, the choice of packages is smaller.

Estimator

The unique feature of SigmaTheta is that it uses the Bayesian statistics (inverse problem) to estimate the variances. The other packages evaluate the *average* value using the classical formulae like (4.4) for AVAR.

Missing data

The Gap Resistant ADEV (GRADEV), included only in AllanTools, enables the evaluation of ADEV in the case of missing data during the measurement.

Error bars

The basic choice is between the classical Lesage-Audoin algorithm [Lesage and Audoin, 1973], and the more sophisticated Greenhall-Riley algorithm [Greenhall and Riley, 2003]. The LA algorithm uses the square root of the 4th moment, assuming that the error distribution is Gaussian centered around the estimate. By contrast, the GR algorithm estimates the degrees of freedom using the χ^2 distribution. For each value of τ , the number of samples and the dominant noise type are taken into account.

However, Wallin points out that the identification of noise processes is also needed [Wallin, 2018]. In the same blog, it is said that AllanTools has no ‘intelligent’ top-level algorithm that takes a time-series as input, and would choose between lag-1-ACF, $B1$ and $R(n)$, and automatically determine the power-law noise type. Wallin suggests that Stable32 may use a combination of lag-1-autocorrelation [Riley and Greenhall., 2004], $B1$ -ratio [Howe et al., 2000, Barnes, 1969], and $R(n)$.

Finally, the principles used inside SigmaTheta are discussed in Sec. 4.2.3.

4.3.2 Software Availability

The availability is summarized in Tab. 4.2. AllanTools and SigmaTheta are open code released to public domain under very similar licenses. Stable32 is © IEEE. The developers informed us⁷ that the code itself is covered by the MIT license, but compiling needs code from other sources. They are trying to clean the code for it to be released into the public domain, provided the IEEE agrees, with the ultimate goal of having at least Linux, macOS and Windows distributions available.

4.4 Stefan Droste

Stefan Droste asked: I guess there are these 4 possibilities:

Pi counter & ADEV

Pi counter & modADEV

⁷Magnus Danielson, private email.

Table 4.2:

SOFTWARE AVAILABILITY AND DEVELOPPERS

AllanTools

Anders E. E. Wallin, Danny Price, Cantwell G. Carson,
Frédéric Meynadier, Yan Xie, and Eric Benkler

Anaconda	https://anaconda.org/conda-forge/allantools
GitHub	https://github.com/aewallin/allantools
PyPI	https://pypi.org/project/AllanTools/

SigmaTheta

François Meyer, François Vernotte, and Attila Kinali, Benoit Dubois

Cmd line	https://gitlab.com/fm-ltfb/SigmaTheta
API	https://gitlab.com/bendub/libsigmatheta Not tested on Mac/Win, some funct. still missing.
GUI	https://gitlab.com/bendub/pysigmatheta β -test on Linux, lives on API via Python wrapper

Stable32

Magnus Danielson, Vivek D. Dwivedi & al.
Originally, William J. Riley (retired)

GitHub	https://github.com/IEEE-UFFC/stable32
IEEE	https://ieee-uffc.org/software
Stable32	http://www.stable32.com/162Stable32.exe

TimeLab

John Miles, KE5FX

Miles Design	http://www.miles.io/timelab/beta.htm
--------------	---

Lambda counter & ADEV

Lambda counter & modADEV

Is there a usecase for all of them or do only some of these make sense?

Chapter 5

Two-Port Devices

ZZZZZZZZZZZZZZZZZZ

5.0.1 Example of Noise in a Two-Port Component

We consider an amplifier having $NF = 2$ dB and flicker coefficient $b_{-1} = 2 \times 10^{-11}$ rad^2 (-107 dBrad²), processing a 10 GHz carrier of power $P = 62.5$ μW (-12 dBm). Let us calculate the noise spectrum and the rms delay fluctuation assuming a low cutoff $f_1 = 10^{-8}$ Hz (reciprocal of 3 years, rounded), and a bandwidth $f_2 = 50$ MHz.

PM noise PSD

Using $b_0 = FkT/P$, in dB we get $10 \log_{10}(b_0) = +2 - 174 + 12 = -160$ dBrad²/Hz, hence $b_0 = 10^{-16}$ rad^2/Hz . The phase noise PSD is

$$S_\varphi(f) = b_0 + \frac{b_{-1}}{f} = 10^{-16} + \frac{2 \times 10^{-11}}{f} \quad \text{rad}^2/\text{Hz}.$$

Using $b_0 = b_{-1}/f$, the corner frequency is $f_c = 2 \times 10^5$ Hz.

Phase-time PSD

Using (6.2) with $\nu_0 = 10$ GHz, we get

$$S_x(f) = k_0 + \frac{k_{-1}}{f} = 2.6 \times 10^{-38} + \frac{5.1 \times 10^{-33}}{f} \quad \text{s}^2/\text{Hz}.$$

RMS delay fluctuation

This is evaluated using (1.32).

White noise $(\delta T)_0 = \sqrt{k_0(f_2 - f_1)} = 1.1 \times 10^{-15}$ s.

Flicker noise $(\delta T)_{-1} = \sqrt{k_{-1} \ln(f_2/f_1)} = 4.3 \times 10^{-16}$ s.

Total $\delta T = \sqrt{(\delta T)_0^2 + (\delta T)_{-1}^2} = 1.2 \times 10^{-15}$ s.

Comment

Common sense suggests that the delay of an amplifier can be of a fraction of a ns (several periods of the 10 GHz carrier) for a wideband device, with a thermal coefficient up to $10^{-3}/K$ if no special design care is taken. As a result, we expect that the thermal effects exceed the random noise, and that the random noise is visible only beyond 10 Hz, where the temperature may be stabilized by the thermal capacitance (inertia).

5.0.2 Suggested Readings About Two-Port Components

Reference [Boudot and Rubiola, 2012] is an extensive treatise of phase noise in amplifiers. Arguably, [Halford et al., 1968] is the first article suggesting that the flicker noise (the parameter b_{-1}) in RF and microwave amplifiers is independent of power and frequency over a rather broad range.

The double-balanced mixer is a tool of paramount importance in PM noise and frequency stability. Reference [Rubiola, 2006] is a useful tutorial, and [Barnes et al., 2011] provides experimental data about noise of commercial double-balanced mixers for phase noise measurements.

We recommend the book [Kester, 2004] to understand analog-digital conversion. Reference [Calosso and Rubiola, 2017] is a tutorial on phase noise digital systems. The origin and the propagation of phase noise in frequency division are explained in [Egan, 1990, Kroupa, 2001].

Phase locking and frequency synthesis are a totally different kind of two-port systems. We suggest [Egan, 2008, Banerjee, 2017, Shu and Sánchez-Sinencio, 2005] for phase locking, [Goldberg, 1999, Kroupa, 1999] for digital frequency synthesis, [Calosso et al., 2012] for the phase noise in digital synthesis, and [Rohde et al., 2021] for a general treatise about modern synthesizers.

Chapter 6

Digital Electronics

This is the arXiv version. There is a more recent version in the directory `Work/2019-Digital/Current`. Figures must be updated too.

This article explains phase noise, jitter, and some slower phenomena in digital integrated circuits, focusing on high-demanding, noise-critical applications. We introduce the concept of *phase type* and *time type* (for short, φ -type and x-type) phase noise. The rules for scaling the noise with frequency are chiefly determined by the spectral properties of these two basic types, by the aliasing phenomenon, and by the input and output circuits.

Then, we discuss the parameter extraction from experimental data and we report on the measured phase noise in some selected devices of different node size and complexity. We observed flicker noise between -80 and -130 dBrad²/Hz at 1 Hz offset, and white noise down to -165 dBrad²/Hz in some fortunate cases and using the appropriate tricks.

It turns out that flicker noise is proportional to the reciprocal of the volume of the transistor. This unpleasant conclusion is supported by a gedanken experiment.

Further experiments provide understanding on: (i) the interplay between noise sources in the internal PLL, often present in FPGAs; (ii) the chattering phenomenon, which consists in multiple bouncing at transitions; and (iii) thermal time constants, and their effect on phase wander and on the Allan variance.

6.1 Introduction

Timing analysis is generally driven by the design of logic functions. That is why specs like “the input must be stable 600 ps before the clock edge” are just countless. From this standpoint, it is sufficient to describe the fluctuations in terms of *jitter*. Broadly speaking, jitter is the time fluctuation, evaluated in reference conditions. Because of the wide bandwidth, jitter is chiefly determined by the white noise. Notice that proper operation requires an analog bandwidth 3–4 times the switching frequency, and in turn up to a few GHz with nowadays components.

When the design comes to spectral analysis and to highly stable oscillators, language and requirements change radically. Fluctuations are generally described in terms of phase noise, expressed either as $S_\varphi(f)$ or $\mathcal{L}(f)$, and the low-frequency phenomena are no longer negligible. Low phase noise is crucial in radars [Skolnik, 2001, Skolnik, 2008, Krieger and Younis, 2006], modern telecomm [Esman et al., 2016], atomic frequency standards [Riehle, 2015] and particle accelerators [Serrano et al., 2011, Jablonski et al., 2015], just to mention some.

In the rapidly changing world of digital electronics, the literature on phase noise is rather old and focuses on frequency dividers, either in TTL and ECL components [Phillips, 1987, Egan, 1990], or in transistor-level modeling. Other references found are more about data transfer in telecom networks than about components [Reinhardt, 2005, Bregni, 2002, Kihara, 1989].

At the time of [Phillips, 1987, Egan, 1990], CMOS technology was used only in microprocessors and complex functions. Gate arrays and FPGAs came later, with a new rapid progress [Mack, 2015, Huang, 2015, Moore’s Law, 2015]. Interestingly for us, gate arrays and FPGAs bridge the gap between logical/computational functions and circuit-level design. The precise control on electrical signals that follows opens a new challenge in understanding noise. However, VLSI engineers are mostly concerned with noise margin, crosstalk, and power distribution [Weste and Harris, 2011]. Conversely, amplitude and phase noise are not studied.

The purpose of this article is to set the basic knowledge about phase noise, and to provide examples. We focus on the clock distribution because clock edges are the most critical ones for timing. This does not sound a limitation, first because critical signals can be synchronized to a clock line, and second because a chip in charge of a highly critical operation should not perform multiple tasks ‘cross-talking’ at random with one another.

Designing the experiments was initially difficult. However, after a noise model and the first results were available, reproducing similar experiments is surprisingly simple. We hope that the reader will be able to port our ideas to other technologies and logic families. The reader may also learn about reverse engineering the noise.

6.2 Definitions, and Phase Noise Models

Phase noise is often expressed as the one-sided PSD $S_\varphi(f)$ of the random phase $\varphi(t)$. In technical literature we often find $\mathcal{L}(f)$, defined as $\mathcal{L}(f) = \frac{1}{2} S_\varphi(f)$ and given in dBc/Hz [?]. Alternatively, phase noise is represented as the phase time fluctuation

Table 6.1: Phase Noise Types and Their Parameters

Noise Type	Dependence on ν_0		Main Equation	Derived Equation	Parameters
	$S_\varphi(f)$	$S_x(f)$			
Pure phase type (pure φ -type)	C	$1/\nu_0^2$	$b_{-1} = \frac{h_{-1}}{V_0^2}$ (6.12)	$k_{-1} = \frac{h_{-1}}{4\pi^2\nu_0^2V_0^2}$ (6.13)	$\sqrt{h_{-1}}$ [V] V_0 [V]
Aliased phase type (aliased φ -type)	$1/\nu_0$	$1/\nu_0^3$	$b_0 = \frac{B h_0}{\nu_0 V_0^2}$ (6.10)	$k_0 = \frac{h_0 B}{4\pi^2\nu_0^3V_0^2}$ (6.11)	$\sqrt{h_0 B}$ [V] V_0 [V]
Pure time type (pure x-type)	ν_0^2	C	$k_{-1} = C$ (6.15)	$b_{-1} = 4\pi^2\nu_0^2 k_{-1}$ (6.16)	$\sqrt{k_{-1}}$ [s]
Aliased time type (aliased x-type)	ν_0	$1/\nu_0$	$k_0 = \frac{J^2}{\nu_0}$ (6.17)	$b_0 = 4\pi^2 J^2 \nu_0$ (6.18)	J [s]

$x(t)$, and its PSD $S_x(f)$. Since $x(t)$ is equivalent to $\varphi(t)$ converted into time, it holds that

$$x(t) = \frac{1}{2\pi\nu_0} \varphi(t) \quad [\text{s}] \quad (6.1)$$

$$S_x(f) = \frac{1}{4\pi^2\nu_0^2} S_\varphi(f) \quad [\text{s}^2/\text{Hz}], \quad (6.2)$$

where ν_0 is the carrier frequency. Our notation is consistent with general literature [?, ?], yet for the choice of fonts for some specific quantities as a minor detail.

A model which is useful to describe phase noise is the polynomial law

$$S_\varphi(f) = \sum_{j=m}^0 \mathbf{b}_j f^j \quad S_x(f) = \sum_{j=m}^0 \mathbf{k}_j f^j, \quad (6.3)$$

where the integer $m < 0$ depends on the device. After (6.2), it holds that $\mathbf{k}_j = \mathbf{b}_j/4\pi^2\nu_0^2$. The sum (6.3) describes the usual noise types: white phase noise \mathbf{b}_0 , flicker phase noise \mathbf{b}_{-1}/f , white frequency noise \mathbf{b}_{-2}/f^2 , etc. Common sense suggests that in two-port components, noise processes higher than $1/f$ (i.e., f^j , $j < -1$) cannot extend over unlimitedly low frequencies, otherwise the input-output delay diverges in the long run.

The polynomial law is also used for the PSD of the voltage noise $n(t)$

$$S_n(f) = \sum_{j=m}^0 \mathbf{h}_j f^j \quad [\text{V}^2/\text{Hz}] \quad (6.4)$$

(notice the font in \mathbf{h}_j , because \mathbf{h}_j reserved for $S_y(f) = \sum_j \mathbf{h}_j f^j$). The reader familiar with analog electronics finds an obvious analogy with the parameter e_n [$\text{nV}/\sqrt{\text{Hz}}$], specified separately for white and flicker noise.

The rms time fluctuation J can be calculated integrating $S_x(f)$ over the system bandwidth (Parseval theorem)

$$J^2 = \int_{f_L}^{f_H} S_x(f) df. \quad (6.5)$$

The lower limit f_L is set by maximum differential delay in the system. The upper limit is $f_H = \nu_0$. The reason is that the fluctuations are sampled at the clock edges, thus at $2\nu_0$. The quantity J^2 can be identified with the variance $\langle x^2(t) \rangle$, yet after filtering out the $f < f_L$ part.

For our purposes, J is approximately equivalent to the *rms jitter*. By contrast, the general term ‘jitter’ has wider scope, mostly oriented to SDH telecomm systems. It includes different types of noise and interferences starting at 10 Hz, with different weight for each (the term ‘wander’ is preferred below 10 Hz). See for example [?, Li, 2008, Reinhardt, 2005] for standards and useful digressions. In a FPGA, there may be a factor 1000 between the rms jitter and the overall jitter, also including interferences.




Figure 6.1: Block diagram describing the noise in the clock distribution.

We introduce two basic types of process discussed below, which take their names from the frequency-scaling properties.

The **phase-type** (or *pure phase-type*) process is, *by definition*, a process in which the statistical properties of $\varphi(t)$ are unaffected after changing the carrier frequency ν_0 in a suitable wide range. Hence, $x(t)$ scales with ν_0 according to (6.1).

The roles of $\varphi(t)$ and $x(t)$ are interchanged in the time-type process. So, the **time-type** (or *pure time-type*) process is, *by definition*, a process in which the statistical properties of $x(t)$ are unaffected after changing the carrier frequency ν_0 in a suitable wide range. Of course, $x(t)$ scales according to (6.1).

The concepts of phase-type and time-type process apply to phase noise, wavelet variances (Allan and Allan-like), environmental effects, etc. Most readers are familiar with the ‘personality’ of the φ -type noise from the phase noise of RF/microwave amplifiers [Boudot and Rubiola, 2012]. Thermal noise, flicker, and some environmental effects in amplifiers behave in this way. Conversely, the thermal drift of the delay in a coaxial cable or optical fiber are time-type processes. The x -type noise also describes the ideal noise-free synthesizer, which transfers $x(t)$ from the input to the output, independently of ν_0 .

6.3 Noise in the Clock Distribution

A lot about phase and time fluctuations can be learned from the simple model sketched in Fig. 6.1. The input signal of frequency ν_0 is first converted into a square wave with full voltage swing, full slew rate and full bandwidth, and then distributed. Restricting our attention to white and flicker, we get the four behaviors listed in Table 6.1 and discussed below.

6.3.1 Spectrum of the Phase-Type (φ -type) Phase Noise

In digital circuits we often encounter the *aliased* φ -type noise. Let us start with φ -type noise at the input of a digital circuit, where the input signal $v(t)$ crosses a threshold affected by a fluctuation $n(t)$. Under the assumption that the input Slew Rate (SR) is high enough to avoid multiple bouncing (Sec. 6.6), we get $x(t) = n(t)/\text{SR}$.

and, after (6.1),

$$\varphi(t) = \frac{2\pi\nu_0}{\text{SR}} n(t). \quad (6.6)$$

Notice that the direct measurement of $n(t)$ is possible only in simple circuits which allow the simultaneous access to input and output of the gate.

The sinusoid is the preferred clock waveform because it propagates through circuit boards with best impedance matching and lowest crosstalk and radiation, and because high purity reference oscillators work in sinusoidal regime. Discarding the dc component and setting the threshold at 0, the clock signal

$$v(t) = V_0 \cos(2\pi\nu_0 t) \quad (6.7)$$

has slew rate $\text{SR}_v = 2\pi\nu_0 V_0$. In this conditions, the phase fluctuation is

$$\varphi(t) = \frac{n(t)}{V_0} \quad (\varphi\text{-type}). \quad (6.8)$$

Generally, the analog bandwidth B of a digital circuit is greater than the max ν_0 by a factor of 3–4. This is necessary for the device to switch correctly. In turn, the bandwidth of $n(t)$ is equal to B . Squaring the input signal samples $n(t)$ at the zero crossings introduces aliasing. The spectrum of the sampled signal is

$$S_{n,s}(f) = \frac{B}{\nu_0} h_0 + \dots \quad (\text{sampled noise}), \quad (6.9)$$

where the $1/f$ and higher terms are neglected because of the comparatively noise power. A trivial way to prove (6.9) is to calculate the variance $\langle n^2(t) \rangle = h_0 B$ (Parseval theorem) before sampling, and to state that it is equal to the variance $\sigma^2 = S_{n,s}(f)\nu_0$ of the sampled signal. Accordingly, the phase noise is

$$b_0 = \frac{h_0 B}{\nu_0 V_0^2} \quad (\text{white, aliased } \varphi\text{-type}) \quad (6.10)$$

$$k_0 = \frac{h_0 B}{4\pi^2 \nu_0^3 V_0^2} \quad (\text{same, after (6.2)}). \quad (6.11)$$

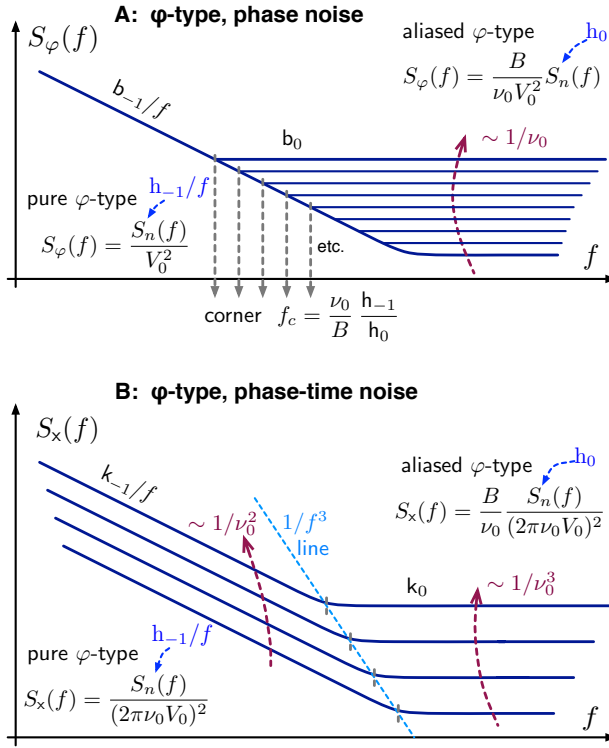
Oppositely, aliasing has negligible effect on flicker h_{-1}/f and on higher terms ($1/f^2$, $1/f^3$ etc.). It follows from (6.8) that

$$b_{-1} = \frac{h_{-1}}{V_0^2}, \quad C \text{ vs. } \nu_0 \quad (\text{flicker, pure } \varphi\text{-type}) \quad (6.12)$$

$$k_{-1} = \frac{h_{-1}}{4\pi^2 \nu_0^2 V_0^2} \quad (\text{same, after (6.2)}). \quad (6.13)$$

Figure 6.2 shows the spectral properties of the φ -type noise. Aliasing scales the white noise as $1/\nu_0$, but it has no effect on flicker. The corner frequency f_c which separates white from flicker regions is obtained equating (6.10) to (6.12)

$$f_c = \frac{\nu_0}{B} \frac{h_{-1}}{h_0} \quad (\text{corner, } \varphi\text{-type noise}). \quad (6.14)$$

Figure 6.2: Spectra originated by the phase type (φ -type) phase noise.

6.3.2 Spectrum of the Time Type (x-type) Phase Noise

The x-type noise originates after the input comparator, where the clock signal has full SR and bandwidth. Though threshold fluctuations are always present, the voltage-to-time conversion has little effect, and the gate is characterized by its delay fluctuations. So, each gate of the clock distribution contributes to the delay, and the fluctuations add up statistically. At a closer sight, the device may be organized hierarchically, for example in gates and cells, likely with a longer propagation time between cells. Nonetheless, the fluctuation is proportional to the length and to the complexity of the distribution chain.

The pure x-type noise is found in the $1/f$ region and below, not affected by aliasing. The noise spectrum is described by

$$k_{-1} = C \text{ vs. } \nu_0 \quad (\text{flicker, pure x-type}) \quad (6.15)$$

$$b_{-1} = 4\pi^2 \nu_0^2 k_{-1} \quad (\text{same, after (6.2)}) \quad (6.16)$$

where k_{-1} is the technical parameter which results from the clock distribution.

The aliased x-type results from sampling the fluctuation at the frequency $2\nu_0$, which affects the white noise region. The spectral parameter k_0 is found in the same




Figure 6.3: Spectra originated by the time type (x-type) phase noise.

way as with (6.9), neglecting the $1/f$ and higher terms

$$k_0 = J^2 / \nu_0 \quad (\text{white, aliased x-type}) \quad (6.17)$$

$$b_0 = 4\pi^2 J^2 \nu_0 \quad (\text{same, after (6.2)}) . \quad (6.18)$$

The spectral properties of the x-type noise — i.e., (6.15)–(6.18) — are summarized in Fig. 6.3. The corner frequency which divides the flicker from the white region is calculated by equating (6.15) to (6.17)

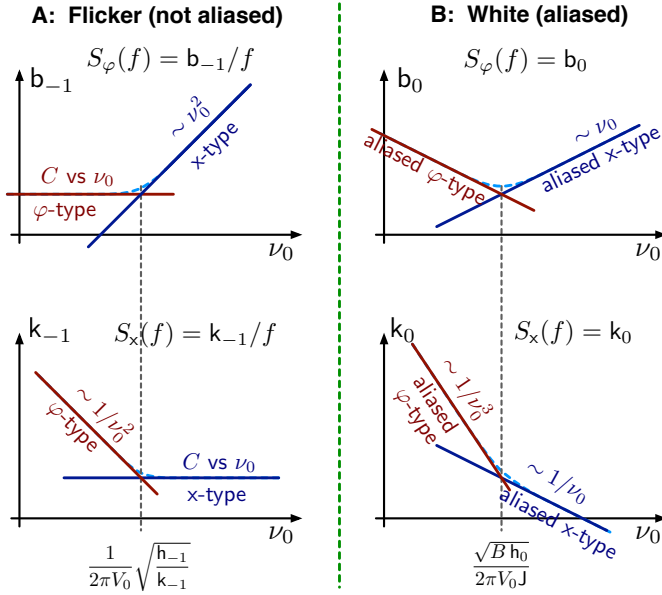
$$f_c = \frac{\nu_0 k_{-1}}{J^2} \quad (\text{corner, x-type noise}) . \quad (6.19)$$

6.3.3 Interpretation of Phase Noise Spectra

A series of spectra $S_\varphi(f)$ taken with several values of ν_0 helps to understand the interplay of noise types. Scaling ν_0 in powers of two seems appropriate.

Let us start with flicker, $S_\varphi(f) = b_{-1}/f$. Comparing (6.12) to (6.16), we expect that the noise is of the φ -type at low ν_0 , and of the x-type at high ν_0 , with a corner frequency

$$\nu_c = \frac{1}{2\pi V_0} \sqrt{\frac{h_{-1}}{k_{-1}}} \quad (\text{flicker}) . \quad (6.20)$$

Figure 6.4: Comparison between φ -type and x-type noise.

This is shown in Fig. 6.4 A. Far from ν_c , we can evaluate

$$h_{-1} = V_0^2 b_{-1} \quad (\nu_0 \ll \nu_c) \quad (6.21)$$

$$k_{-1} = \frac{b_{-1}}{4\pi^2 \nu_0^2} \quad (\nu_0 \gg \nu_c). \quad (6.22)$$

The white phase noise $S_\varphi(f) = b_0$ is described by (6.10) at low ν_0 , and by (6.18) at high ν_0 , separated by the cutoff

$$\nu_c = \frac{\sqrt{B h_0}}{2\pi V_0 J}. \quad (\text{white}). \quad (6.23)$$

This is shown on Fig. 6.4 B. At low ν_0 , (6.10) enables to calculate the noise power $\langle n^2(t) \rangle = h_0 B$ of the input threshold

$$h_0 B = V_0^2 b_0 \nu_0 \quad (\nu_0 \ll \nu_c). \quad (6.24)$$

Assuming that B is equal to 3–4 times the maximum ν_0 , we can infer h_0 and the noise voltage $e_n = \sqrt{h_0}$. Conversely, at high ν_0 we can extract the fluctuation

$$J = \frac{1}{2\pi} \sqrt{\frac{b_0}{\nu_0}} \quad (\nu_0 \gg \nu_c). \quad (6.25)$$

This can be compared to the rms jitter, if available in the specs.




Figure 6.5: The digital phase meter is either a Symmetricom (now Microsemi) 5125 or 5120. The two outputs may have different frequency.

6.4 Selected Noise Measurements

We measured the phase noise of several devices routinely used in our labs. This is a necessary step, before considering an unbound search for the best. Accordingly, the measurement method (Fig. 6.5) is more about flexibility than about sensitivity. Anyway, the phase noise of digital components is generally higher than that of common low noise components (i.e., amplifiers and mixers). On the other hand, we need simple operation in a wide range of frequency, with signals that may *not* be at the same frequency as the reference. For us, this is the relevant feature of the Microsemi 5125 (1–400 MHz) and 5120 (1–30 MHz) instruments. These instruments make use of correlation and average on the spectra of two nominally equal channels which measure the same quantity, which rejects the single channel noise [?, ?]. Notice that the oscillator is common mode, with very small differential delay, hence its noise is highly rejected. The Fourier frequency spans from 1 mHz to 1 MHz.

6.4.1 Cyclone III (65 nm)

In a first experiment, we measure a Cyclone III [?] in a clock buffer configuration. The input sinusoidal clock $V_0 = 1 V_{\text{peak}}$ (+10 dBm on 50 Ω) is squared and distributed as in Fig. 6.1 A. The spectrum is shown in Fig. 6.6.

We first look at the white noise region. Our model suggests aliased φ -type noise (6.10) at low ν_0 , and aliased x-type noise (6.18) beyond the cutoff given by (6.23), as shown on Fig. 6.4 B. Starting from $\nu_0 = 3.125$ MHz, b_0 scales down as -3.5 dB per factor-of-two, in fairly good agreement with the 3 dB predicted by the model. This results from the data fit shown on Fig. 6.6 top-right. Taking $V_0 = 1$ V, (6.10) gives a threshold fluctuation $\sqrt{h_0 B} = 550 \pm 65 \mu\text{V}$. The $\pm 65 \mu\text{V}$ results from $b_0 \propto 1/\nu_0^{1.16}$, instead of the $1/\nu_0$ law. Assuming $B = 2.5$ GHz (analog bandwidth, four times the maximum toggling frequency), we get $\sqrt{h_0} = 11 \pm 1.3 \text{ nV}/\sqrt{\text{Hz}}$. This is in agreement with general experience, which suggests that general high-speed electronics has a typical noise level of 10–15 nV/ $\sqrt{\text{Hz}}$.

At $\nu_0 \geq 100$ MHz, the white noise falls outside the 1 MHz span. Since this occultation occurs before the aliased x-type noise shows up, we have no direct access to k_0 . On Fig. 6.6, at the maximum f (1 MHz) and at 400 MHz carrier, the white noise is below $-138 \text{ dBrad}^2/\text{Hz}$ (upper bound). This value, integrated over $B = 400$




Figure 6.6: Phase noise of the Cyclone III clock distribution.

MHz and converted into time, gives 1 ps, which is an upper bound for J .

Flicker noise is in good agreement with pattern of Fig. 6.4 B only at $\nu_0 \geq 100$ MHz. From this part of the plot, we calculate $\sqrt{k_{-1}} = 21$ fs. By contrast, at $\nu_0 \leq 50$ MHz b_{-1} scales as ≈ 1.5 dB per factor-of-two instead of being constant. This discrepancy is not understood. However, the $1/f$ region is rather irregular, and corrupted by bumps, even more pronounced at low ν_0 .

The lowest flicker found on Fig. 6.6 (-115 dBrad²/Hz at 3.125 MHz carrier), converted into voltage using (6.12), gives $\sqrt{h_{-1}} = 2.6$ μ V (upper bound for the input voltage flicker). Interestingly, this value is similar to the flicker of some CMOS high-speed operational amplifiers (for instance, 1.9 μ V for the Texas Instruments OPA354A).

Figure 6.7 shows the phase noise of the output buffer. The white noise is too low to be visible with the 1 MHz span, masked by flicker and by some bumps at $10^4 \dots 10^6$ Hz. By contrast, the flicker noise is in perfect agreement with the 6 dB per factor-of-two model (pure x -type noise). Comparing Fig. 6.7 to Fig 6.6, at $\nu_0 = 400$ MHz

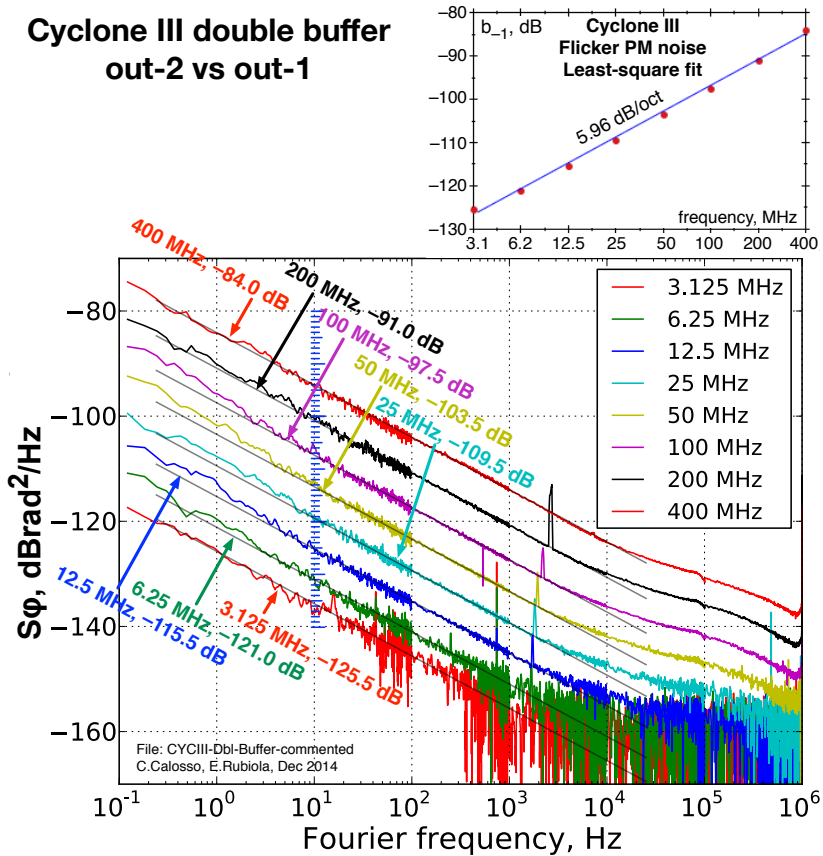


Figure 6.7: Phase noise of the Cyclone III, measured by comparing two outputs. Take away 3 dB for the noise of one buffer.

the flicker of the complete clock distribution is close to that of the output buffer. So, the contribution of the output buffer is not negligible. Conversely, at lower ν_0 a significantly larger flicker rises in the clock distribution chain.

6.4.2 Measuring the Time Type (x-Type) Noise with the Λ Divider

After some tests, we realized that the Λ frequency divider [?] is a good tool to measure the x-type noise of the clock distribution. First, a frequency divider is useful in that the input time fluctuation (φ -type noise, (6.13)) is kept low by using a high input frequency, while the measurement at the lower output frequency is simpler (both instruments are suitable, and the background is lower). Second, the Λ divider circumvents the aliasing phenomenon. In fact, a Λ divider $\div \mathcal{D}$ provides a triangle-like output waveform by combining \mathcal{D} phases of a square wave, which is equivalent to sampling at the input frequency.




Figure 6.8: Phase noise of some components used as a $\div 10$ frequency divider in the Λ configuration.

Figure 6.8 shows the phase noise of some devices used as $\div 10$ dividers in Λ configuration, with 100 MHz input and 10 MHz output frequency. The flicker coefficient is clearly identified, not corrupted by artifacts. The bump at 20 kHz (Zynq and Cyclone III) is due to the insufficiently filtered power supply. Finally, the Λ divider implemented with the Max 3000 deserves mentioning for its low noise ($b_{-1} = -130.5$ dBrad and $b_0 = -165$ dBrad²/Hz). This is lower than regular dividers (general experience), and just 10 dB above the NIST regenerative dividers [Hati et al., 2013] at the same output frequency.

6.5 The Volume Law

The idea that the phase noise coefficient b_{-1} is proportional to $1/V$, where V is the active volume, has been around for a while. In quartz resonators, this appears either directly or as a side effect of the larger size at lower frequency [Kroupa, 1998, Kroupa, 2005, ?, Walls et al., 1992, Driscoll and Hanson, 1993, Sthal et al., 2013]. In ultrastable Fabry-Perot cavities, flicker is powered by thermal noise and proportional to the reciprocal of the length [Saulson, 1990, Numata et al., 2004] which is approximately equivalent to $1/V$ after mechanical design rules.

The $1/V$ law results from a gedankenexperiment in which we combine m equal and independent devices, giving $b_{-1}|_{\text{total}} = b_{-1}|_{\text{dev}}/m$. This has been confirmed




Figure 6.9: Flicker coefficient b_{-1} of digital devices, related to the cell size \mathcal{S} .

experimentally with amplifiers [?, Chapter 2], [Boudot and Rubiola, 2012]. Flicker is of microscopic origin because the probability density function is Gaussian, which originates from a large statistically-independent population through the central limit theorem. So, the m devices can be combined in a factor-of- m larger device exhibiting a factor-of- $1/m$ lower flicker. Similarly, we expect higher flicker if the size of the device is scaled down, until space correlation appears. The limit for small volume is not known.

In digital electronics, the volume V of the active region is proportional to the node size \mathcal{S} . For reference, \mathcal{S} is of $10\ \mu\text{m}$ in Intel 4004 (1971), and of $16\ \text{nm}$ in the Apple A10 Fusion chip of the iPhone 7. While the footprint surface is proportional to \mathcal{S}^2 , the two scaling rules are common in the literature on VLSI systems, known as *constant-voltage* and *Dennard* [Weste and Harris, 2011, P. 253], [Dennard et al., 1974], agree in the depth proportional to \mathcal{S} . Thus, $V \propto \mathcal{S}^3$. The wire delay may contain $\sqrt{\mathcal{S}}$, however, the flicker associated to wires is too small to deserve attention [Verbruggen et al., 1989].

We measured a few components using the $\div 10\ \text{A}$ divider configuration. This gives access to the $1/f$ noise of the clock distribution, which is of the x-type. We used $100 \rightarrow 10\ \text{MHz}$, or $30 \rightarrow 3\ \text{MHz}$ with the Cyclone and the Cyclone II for practical reasons, sharing a 5125A and a 5120A. The results are shown in Fig. 6.9, which compares the $1/f$ PM noise to \mathcal{S} .

The MAX V is not accounted for in the analysis because the spectrum was taken in unfavorable conditions, yet kept for completeness. A linear regression gives $k_{-1} = -26.2 \log_{10}(\mathcal{S}) - 219.5\ \text{dBs}^2$, with \mathcal{S} in nm. Fitting the same data with the exact volume law gives $k_{-1} = -30 \log_{10}(\mathcal{S}) - 212.1\ \text{dBrad}^2/\text{Hz}$. The $-26.2\ \text{dB/dec}$ slope is




Figure 6.10: Simulation of carrier crossing a fluctuating threshold (normalized 1 Hz carrier, 1 V_{peak}). Multiple crossing occurs in the center of the plot.

reasonably close to the $1/V$ law (-30 dB/dec), with a number of measurement and accuracy insufficient to assess a discrepancy.

6.6 Input Chatter

Chatter is a fast random switching of a comparator, which occurs in the presence of wideband noise when the mean square slew rate of noise exceeds that of the signal at the threshold, i.e., $\langle \text{SR}_n^2 \rangle > \text{SR}_v^2$. The phenomenon is shown in Fig. 6.10 and 6.11.

Following the Rice's approach [?, ?], noise in the small interval $[f, f + \Delta f]$ can be represented as the sinusoidal signal $n_f(t) = V_f \cos(2\pi f t + \theta_f)$, which has random amplitude V_f , random phase θ_f , and slew rate

$$\text{SR}_{n,f} = 2\pi f V_f \sin(\theta_f). \quad (6.26)$$

The Parseval theorem requires that $\langle n_f^2(t) \rangle = S_n(f) \Delta f$, thus

$$\langle V_f^2 \rangle = 2S_n(f) \Delta f \quad (6.27)$$

because $\langle \cos^2(\dots) \rangle = 1/2$ in $n_f(t)$. The mean square slew rate is calculated combining (6.26) and (6.27), integrating on frequency, and averaging on θ_f . Since $\langle \sin^2(\theta_f) \rangle = 1/2$,

$$\langle \text{SR}_n^2 \rangle = 4\pi^2 \int_0^\infty f^2 S_n(f) df. \quad (6.28)$$

In turn, $\langle \text{SR}_n^2 \rangle$ is determined by white noise $S_n(f) = h_0$, $f = [0, B]$. Other noise types are negligible because they occur at low frequency, compared to B , and because




Figure 6.11: Example of chatter (multiple bouncing) when the input SR is insufficient as compared to the SR associated to noise.

of the f^2 term in (6.28). Thus

$$\langle \text{SR}_n^2 \rangle = \frac{4\pi^2}{3} h_0 B^3. \quad (6.29)$$

Since the clock signal (6.7) has slew rate $\text{SR}_v = 2\pi\nu_0 V_0$, the chatter threshold is

$$\nu_0 V_0 = \sqrt{\frac{1}{3} h_0 B^3} \quad (\text{chatter threshold}). \quad (6.30)$$

Taking the Cyclone III parameters (Sec. 6.4.1, $B = 2.5$ GHz and $e_n = 11$ nV/ $\sqrt{\text{Hz}}$, thus $h_0 = 1.21 \times 10^{-16}$ V²/Hz), and $\nu_0 = 4.7$ MHz, (6.30) suggests a threshold $V_0 = 169$ mV. On Fig. 6.11, we see that chattering occurs at $V_0 = 100$ mV, and at $V_0 = 50$ mV the transitions are broken. Given the difficulty of identifying the parameters, the agreement between model and observation is satisfactory.

After (6.30), chattering is more likely at low carrier frequency. However, Fig. 6.11 shows that this can occur at 5 MHz, a standard frequency of great interest for high stability signals.

6.7 Internal PLL

The internal PLL is intended to provide high frequency internal clock stabilized to an external reference, often 5-10-100 MHz. We show simple experiments which give

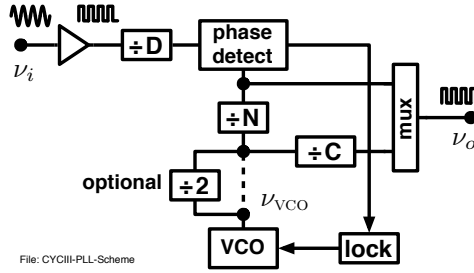


Figure 6.12: Cyclone III internal PLL frequency multiplier.

insight in the Cyclone III.

The PLL is shown in Fig. 6.12. The VCO operates in the 0.6–1.3 GHz range, extended to 300–650 MHz by the optional $\div 2$ divider, always present in our tests. A classical phase-frequency detector (PFD) is present, with charge pump output driving the analog feedback to the VCO. The PLL output frequency is $\nu_o = \frac{N}{CD}\nu_i$. This leaves three degrees of freedom (N , C and D), two of which are available to the designer. The programming tool (Quartus) uses one to ensure that internal design rules are satisfied.

The VCO relies on a LC resonator on chip. General literature suggests a quality factor Q of 5–10, limited by the technology [Hajimiri and Lee, 1999]. Therefore, we expect a Leeson frequency $f_L = \nu_{VCO}/2Q$ of the order of 50 MHz.

In a first experiment (Fig. 6.13), we use the PLL as a ‘cleanup’ ($\nu_o = \nu_i$), yet with a high purity input. This gives the noise of the PLL, at different values of ν_i . For lowest noise, we use the phase comparator at the highest possible frequency (ν_i) by setting $D = 1$. The VCO frequency ends up to be 400, 600 or 640 MHz, depending on ν_o . On Fig. 6.13, the white noise floor is not seen. This is sound because noise can be white only beyond f_L , which is beyond the 1 MHz span. Flicker is of the φ -type at 5 and 10 MHz, with $b_{-1} = 2.5 \times 10^{-10}$ rad²/Hz (−96 dB). Since this type of noise is *not* scaled down by the $\div N$ divider in the loop, we ascribe it to the *phase detector*. This is because (i) with the tight lock implemented we do not expect to see the VCO; and (ii) the input comparator and the output stage of the $\div N$ divider have some 10 dB lower noise in similar conditions (−115 dBrad²/Hz, Section 6.4.1).

In the second experiment, we use the PLL as a frequency multiplier in powers of two ($\nu_o = 2^m \nu_i$) from 10 MHz to 640 MHz, with $\nu_i = 10$ MHz. Again, we use $D = 1$ for lowest noise. The VCO delivers 320, 400 or 640 MHz, depending on ν_o . The phase noise spectrum (Fig. 6.14) indicates that flicker is of the x -type, scaling up as ν_o^2 . This indicates that the phase detector is the dominant source of noise, with negligible contribution of the dividers. So, the time fluctuation $x(t)$ is transferred from the phase detector to the VCO, and then from the VCO to the output. The phase $\varphi(t)$ scales accordingly, that is, $\propto N/C$.




Figure 6.13: The internal PLL is used as a buffer, that is, $\nu_o = \nu_i$.

6.8 Thermal Effects

6.8.1 Thermal Transients

Common sense suggests that delay is affected by the junction temperature T_J , while other parameters like T_C and T_A (case and ambient temperature) are comparatively smaller importance.

Our method consists in using the electrical power P to heat the chip, and calculate T_J from the thermal resistance Θ_{JA} and the transients. In turn, P is chiefly set by the charge/discharge cycle of the gate capacitance, whose energy is $E = CV^2$. Thus, N gates switching at ν_0 dissipate $P = NCV^2\nu_0$. Of course, P can be changed instantaneously. The delay is measured with a Symmetricom 5125A test set used as a phase meter and also as a time-interval counter.

We measured a Cyclone III used as a clock buffer (actually, 10 buffers connected in parallel through 330 Ω resistors). The temperature had to be low-pass filtered by covering the card with a small piece of tissue. The results are shown in Fig. 6.15.

In the main body, all the curves show an exponential behavior plus a linear drift

$$x(t) = k' \Delta T (1 - e^{-t/\tilde{\tau}}) + k'' t, \quad (6.31)$$

where $\Delta T = T_J - T_A$ results from setting ν_0 in powers of two, and $\tilde{\tau}$ is the time constant. For reference, we observed $P = 1$ W at 400 MHz, which means $\Delta T \approx 10$ K with $\Theta_{JA} \approx 10$ K/W (including the thermal pad on the pcb), and neglecting the dissipation at $\nu_0 = 0$.




Figure 6.14: The internal PLL is used as a frequency multiplier in powers-of-two of multiples of the 10 MHz frequency reference.

The linear drift (1 fs/s, or 10^{-15} fractional frequency) does not scale with power. This behavior is typical of the environment temperature, slowly drifting during the measurement (a fraction of a Kelvin over 1 hour). Extrapolating the drift to $t = 0$, we get the asymptotic effect of the ΔP transient alone.

The time constant $\tilde{\tau}$ is found as the intercept of the tangent at $t = 0$ and the linear drift (dashed lines). This graphical process removes the drift. The value $\tilde{\tau} = 400$ s is the same for all the transients.

The inset of Fig. 6.15 shows the delay versus the carrier frequency (dissipated power). As expected, the delay is proportional to T_J , set through ν_0 . Accounting for P and Θ_{JA} , the thermal coefficient of the delay is 10 ps/K.

6.8.2 Allan Deviation

Generally, $\sigma_y(\tau)$ should follow the $1/\tau$ law (white and $1/f$ phase noise). Other types of instability, as frequency noise would reveal a phase noise steeper than $1/f$, and the delay of the device would diverge in the long run. However, bumps may be present. Notice that $1/f$ phase noise in practice never yields large integrated delay.

Figure 6.16 shows the Cyclone III Allan deviation $\sigma_y(\tau)$, measured with a Symmetricom 5125A test set.

We first discuss the $1/\tau$ region of Fig. 6.16 A. At low ν_0 , $\sigma_y(\tau)$ decreases proportionally to $1/\nu_0$. For $\tau = 1$ s, we read $\sigma_y = 10^{-12}$ at 3.125 MHz, 5×10^{-13} at 6.25




Figure 6.15: Thermal effects measured on a Cyclone III FPGA. Each curve represents the thermal transient when the clock frequency is divided by two.

MHz, etc. At higher ν_0 the curves get closer to one another, and overlap at $\nu_0 \geq 100$ MHz.

Taking the classical conversion formulae for Allan variance and spectra (for example, [Kroupa, 1983, P. 77–80], or [?]), the $1/\nu_0$ behavior is equivalent to $h_1 \propto 1/\nu_0^2$ (frequency fluctuation spectrum $S_y(f) = h_1 f$), thus to $b_{-1} = C$ vs. ν_0 . This is the signature of the pure φ -type noise, as expected at low ν_0 and at low f , thus at long τ . We recall that the fluctuation of the input threshold is dominant at low ν_0 , and that the low f region is dominated by the $1/f$ phase noise, virtually unaffected by aliasing.

By contrast, the $\sigma_y(\tau) = C$ vs. ν_0 behavior is equivalent to $h_1 = C$ vs. ν_0^2 , thus $b_{-1} \sim \nu_0^2$. This is the typical of the pure x-type noise, as expected at high ν_0 and at low f , thus at long τ . The fluctuation of the input threshold is no longer relevant, and the low f region is still dominated by the $1/f$ phase noise, virtually unaffected by aliasing.

In summary, the $1/\tau$ region of the $\sigma_y(\tau)$ plot is consistent with the predictions of Section 6.2.

On the right hand of Fig. 6.16 A, $\sigma_y(\tau)$ seems to leave the $1/\tau$ law. This can only be a local phenomenon, i.e. a bump. Carrying on the experiment, in Fig. 6.16 A the measurement of $\sigma_y(\tau)$ restarts immediately after switching ν_0 , while in Fig. 6.16 B the

Cyclone III clock buffer

File: Cyclone-III-adev-VS-Idle-time-H
C.Calosso, E.Rubiola, Aug 2014




Figure 6.16: Allan deviation $\sigma_y(\tau)$ derived from the FPGA delay.

measurement of $\sigma_y(\tau)$ is delayed by 1 hour after switching ν_0 . The relevant difference is that in A each curve suffers from the cooling-down transient of the previous measurement, while in B each measurement starts in steady state. Bumps show up in A at $\tau \geq 30$ s, and they get stronger at higher ν_0 , where the thermal dissipation is stronger, and almost disappear in B. This is a qualitative confirmation of the presence of two separate time constants (end of Sec. 6.8.1).

6.8.3 Side Effects of the Thermal Dissipation

We have shown that the electrical activity inside the FPGA heats the chip, and in turn affects the delay. Variations exceeding 50 ps have been observed in the presence of a light burden. The analysis gives a warning, thermal crosstalk is around the corner when the same FPGA is in charge of more than one task, made worse by the heat latency. Attempts to fit low noise and high-stability functions (frequency dividers, etc.) in a chip processing at high rate may be difficult or give unpredictable results.

Acknowledgments

This work is a part of the “Programme d’Investissement d’Avenir” projects in progress in Besancon, i.e., Oscillator IMP, First-TF, and Refimeve+. Funds come from the

ANR, the Region Franche Comté, INRIM, and EMRP Project IND 55 Mclocks.

We thank the Go Digital Working Group for general help and fruitful discussion, and among them chiefly Jean-Michel Friedt, Pierre-Yves “PYB” Bourgeois, and Gwenhael “Gwen” Goavec-Mérou.

Chapter 7

Systems

Page XXIII-4

$$\ddot{x} + \frac{\omega_n}{Q} \dot{x} + \omega_n^2 x = \frac{\omega_n}{Q} \dot{v} \quad (7.1)$$

$$s^2 + \frac{\omega_n}{Q} s + \omega_n^2 = 0 \quad (7.2)$$

$$s = \sigma_p + i\omega_p \quad (7.3)$$

$$\sigma_p = -\frac{\omega_n}{2Q} \quad (7.4)$$

$$\omega_p = \frac{\omega_n}{2Q} \sqrt{4Q^2 - 1} \quad (7.5)$$

$$\tau = \frac{2Q}{\omega_n} \quad (7.6)$$

$$\beta(s) = \frac{\omega_n}{2Q} \frac{s}{s^2 + \frac{\omega_n}{2Q} s + \omega_n^2} \quad (7.7)$$

$$\Omega = \omega_0 - \omega_p \quad \text{thus} \quad \omega_p = \omega_0 - \Omega \quad (7.8)$$

$$x_{\text{off}} = \cos(\Omega_p t) e^{-t/\tau} \cos(\omega_0 t) + \sin(\Omega t) e^{-t/\tau} \sin(\omega_0 t) \quad (7.9)$$

$$\begin{aligned} x_{\text{on}} &= \mathcal{A} \cos(\omega_p t) e^{-t/\tau} + \mathcal{B} \sin(\omega_p t) e^{-t/\tau} + \mathcal{C} \cos(\omega_0 t) + \mathcal{D} \sin(\omega_0 t) \\ &= \left[\mathcal{C} + \mathcal{A} \cos(\Omega t) e^{-t/\tau} - \mathcal{B} \sin(\Omega t) e^{-t/\tau} \right] \cos(\omega_0 t) + \\ &\quad + \left[\mathcal{D} + \mathcal{A} \sin(\Omega t) e^{-t/\tau} + \mathcal{B} \cos(\Omega t) e^{-t/\tau} \right] \sin(\omega_0 t) \end{aligned} \quad (7.10)$$

Phase step κ . Boundary conditions

$$\begin{aligned} t \rightarrow \infty &\Rightarrow e^{-t/\tau} = 0 \\ &\Rightarrow \mathcal{C} = 1, \quad \mathcal{D} = -\kappa \end{aligned} \quad (7.11)$$

$$\begin{aligned} t \rightarrow 0 &\Rightarrow e^{-t/\tau} = 1 \\ &\Rightarrow \mathcal{C} + \mathcal{A} = 0 \rightarrow \mathcal{A} = -1 \end{aligned} \quad (7.12)$$

$$\Rightarrow \mathcal{D} + \mathcal{B} = 0 \rightarrow \mathcal{B} = -\kappa \quad (7.13)$$

Chapter 8

Power Splitters¹

8.1 Basic facts

Modern phase noise analyzers use two separate and independent channels which measure simultaneously the same device under test, as shown in Fig. 8.1. The noise of the phase detector and of the references is rejected after correlating and averaging on a sufficiently large number m of measures. Such approach requires a power splitter at the analyzer's input. The power splitter found in phase-noise analyzers are shown in Fig. 8.2, but more exotic configurations are possible. The most commonly used splitter is a 3 dB directional coupler terminated at one input, either of the 180° type (shown), or of the equivalent 90° type. Denoting with P_0 the available power² at the input, the available power at each output is $P_0/2$. The Y splitter, found in Keysight instruments dissipates $P_0/2$ in the internal resistors, thus the available power at each output is $P_0/4$. The V splitter, found in Holzworth instruments, has the same loss of the Y type and some advantages³, at the cost of poorer impedance matching.

The naive representation of Fig. 8.1 suggests that the DUT phase noise is replicated identically at two outputs of the splitter, thus the noise of the detectors and of the references can be averaged out by an appropriate statistical algorithm. A deeper analysis (Fig. 8.3) reveals the following problems

¹A significant part of the material in this chapter comes from Y. Gruson et al., “Cross-Spectrum PM Noise Measurement...”, IEEE T UFFC 2017, open access.

²The available power is the maximum delivered power, which occurs in conjugate matching conditions.

³This is explained by Jason Breitbath in a conference article.





Figure 8.1: .

3 dB directional coupler



resistive power splitters




Figure 8.2: .

1. The splitter loss deteriorates the signal-to-noise ratio (SNR). The consequence is that a longer averaging time is needed for the measure to converge to the DUT noise, but the asymptotic result is free from errors.
2. The power splitter introduces anti-correlated noise originated by the thermal energy associated to the internal dissipation. This is self-evident in the 3 dB splitter, where the thermal noise of the dark-port resistor goes to the two outputs with equal and opposite coefficients. Something similar happens in the resistive schemes.

The consequence is that the instrument under-estimates the DUT noise. Fortunately, such error can be modeled and corrected, leaving a comparatively small uncertainty.

3. If the two outputs of the splitter are not perfectly isolated from one another, the noise back-scattered from the input of one phase detector shows up at the input of the other detector. This originates correlated noise, and in turn a systematic error in the measurement of the DUT noise. Sadly, this effect is highly dependent on the implementation. Modeling it, or even assessing the upper bound in a wide range of carrier power and frequency is a challenge.

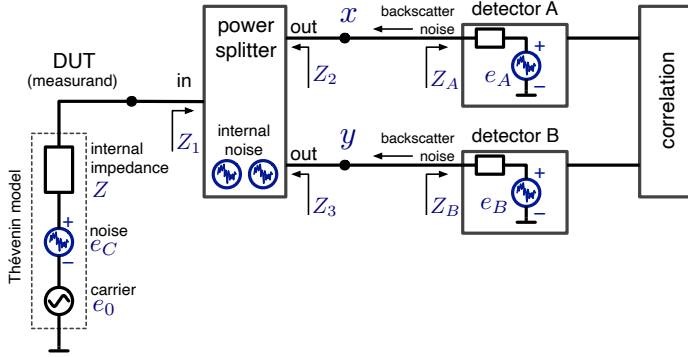


Figure 8.3: Two-channel scheme. (Replace $e_0 \rightarrow V_0$ for consistency to the clock signal, and also to avoid confusion between rms and peak values).

Notice that from our standpoint e_C is the *signal*, and all the other random generators are the *noise* (background noise).

The following Sections provide a basic analysis of errors and uncertainties introduced by the power splitter.

8.1.1 Thévenin model of carrier and noise

A noise-free oscillator is modeled as a sinusoidal generator of peak amplitude V_0 and frequency ν_0 , in series to a resistor R_0 , equal to the characteristic impedance of the system. Accordingly, the available power is

$$P_0 = \frac{V_0^2}{8R_0} \quad [\text{W}] \quad (8.1)$$

The thermal noise of a resistor R at the temperature T is generally modeled as a random generator of rms emf⁴ e ruled by

$$e^2 = 4kTR \quad [\text{V}^2] \text{ in 1 Hz bandwidth} \quad (8.2)$$

where $k = 1.38 \times 10^{-23} \text{ J/K}$ is the Boltzmann constant. This is the same as stating that the PSD of e is

$$S_e = 4kTR \quad \text{V}^2/\text{Hz} \quad (8.3)$$

When the generator is impedance matched to the load, the available voltage and power are

$$e_0^2 = kTR \quad [\text{V}^2/\text{Hz}] \quad (8.4)$$

$$S = kT \quad [\text{W}/\text{Hz}] \quad (8.5)$$

⁴Electromotive force




Figure 8.4

Referring to Fig. 8.3, ‘ e ’ is e_C for the oscillator noise, and e_A and e_B for the noise back-scattered from the detectors. The subscripts may be adapted to the context.

Combining carrier and thermal noise, the phase noise is

$$S_\varphi = \frac{e_C^2}{V_0^2/2} = \frac{kT}{P_0} \quad (8.6)$$

By analogy with thermal noise, an equivalent temperature can be associated to the oscillator’s phase noise

$$T_{\text{eq}}(f) = \frac{P_0 S_\varphi(f)}{k} \quad (8.7)$$

We agree that the phase noise more complex than the additive of the above formula, but this representation is satisfactory for the treatise below.

Now we review the passive power splitters commonly used in phase noise analyzers, focusing on the systematic errors and uncertainty. Reference [Nelson]⁵ provides an independent analysis, including more exotic power splitters.

8.1.2 Noise theory of the dissipative attenuator (refer to somewhere else)

8.2 Common power splitters, impedance matched




Figure 8.5: Directional coupler used as the 3-dB power splitter, and noise emfs. (Replace $e_0 \rightarrow V_0$ for consistency to the clock signal, and also to avoid confusion between rms and peak values).

8.2.1 Directional coupler

The 3 dB power splitter is a loss-free directional coupler where the unused input port⁶ is terminated. Figure 8.5 shows the 3 dB power splitter connected to the oscillator under test (DUT), together with the noise sources. In impedance matching conditions, the output fluctuations are

$$x = \frac{1}{2}e_A + \frac{1}{2\sqrt{2}}e_C - \frac{1}{2\sqrt{2}}e_D \quad (8.8)$$

$$y = \frac{1}{2}e_B + \frac{1}{2\sqrt{2}}e_C + \frac{1}{2\sqrt{2}}e_D \quad (8.9)$$

The cross spectrum of the fluctuations is

$$\begin{aligned} S_{yx} &= \frac{1}{8} [e_C^2 - e_D^2] \\ &= \frac{1}{2}k [T_C - T_D] R_0. \end{aligned} \quad (8.10)$$

Equation (8.10) is the physical principle of the correlation radiometer⁷, and also used in Johnson thermometry⁸. In this case there is no carrier, T_D is the reference temperature, chiefly the triple point of H_2O , and the instrument measures the difference $\Delta T = T_C - T_D$. The preferred estimator⁹ is

$$\widehat{\Delta T} = \frac{2}{kR_0} \langle \Re \{ S_{yx} \} \rangle_m \quad (\text{thermometer}). \quad (8.11)$$

⁵An article by Craig W. Nelson analyzes more exotic power splitters.

⁶For a loss-free passive reciprocal device to be impedance matched at all ports, it must have an even number of ports. So, a two-way power splitter has an additional input, called *dark port* terminated to an internal resistor.

⁷Allred-1962

⁸White-1996-Thermometry

⁹The reason for $\Re \{ S_{yx} \}$ instead of S_{yx} is explained somewhere else.

For application to the measurement of phase noise, we replace

$$\begin{aligned} e_A^2 = e_B^2 &\rightarrow 4kT_{\text{bs}} && \text{(back-scatter equivalent temperature)} \\ e_C^2 &\rightarrow 4kT_{\text{eq}} && \text{(oscillator's equivalent temperature)} \\ e_D^2 &\rightarrow 4kT_s && \text{(splitter's dark-port temperature)} \end{aligned}$$

thus

$$S_{yx} = \frac{1}{2} k (T_{\text{eq}} - T_s) \quad (8.12)$$

Reminding that the available power is $P_0/2$ at each output, a naive choice for the estimator of the phase noise is

$$\widehat{S}_\varphi = \frac{\Re\{S_{yx}\}}{P_0} \quad (8.13)$$

which gives

$$\widehat{S}_\varphi = \frac{k(T_{\text{eq}} - T_s)}{P_0} \quad \text{(affected by systematic error)} \quad (8.14)$$

This estimator is affected by a systematic error

$$\Delta S_\varphi = \widehat{S}_\varphi - S_\varphi \quad (8.15)$$

which is equal to

$$\Delta S_\varphi = -\frac{kT_s}{P_0} \quad \text{(systematic error)} \quad (8.16)$$

with respect to the true $S_\varphi = kT_{\text{eq}}/P_0$. The systematic error is fixed by replacing the estimator with

$$\widehat{S}_\varphi = \frac{\Re\{S_{yx}\} + kT_s}{P_0} \quad \text{(estimator corrected for } T_s) \quad (8.17)$$

8.2.2 Directional coupler followed by 3 dB attenuators




Figure 8.6: Resistive (Y) 6-dB power splitter and noise emfs. **Warning, I changed the sign of the splitter's internal generators. (Replace $e_0 \rightarrow V_0$ for consistency to the clock signal, and also to avoid confusion between rms and peak values).**

8.2.3 Y resistive power splitter

The Y resistive network (Fig. 8.6) is sometimes used instead as a power splitter. To our knowledge, the Keysight E5500 series¹⁰ is the one and only case where a commercial instrument uses the Y coupler. Dropping the carrier, the random signals at the splitter output are

$$x = \frac{1}{2}e_A + \frac{1}{4}e_B + \frac{1}{4}e_C + \frac{1}{2}e_a - \frac{1}{4}e_b - \frac{1}{4}e_c \quad (8.18)$$

$$y = \frac{1}{4}e_A + \frac{1}{2}e_B + \frac{1}{4}e_C - \frac{1}{4}e_a - \frac{1}{2}e_b - \frac{1}{4}e_c \quad (8.19)$$

$$xy = \frac{1}{8}e_A^2 + \frac{1}{8}e_B^2 + \frac{1}{16}e_C^2 - \frac{1}{8}e_a^2 - \frac{1}{8}e_b^2 - \frac{1}{8}e_c^2 \quad (8.20)$$

Application to the measurement of phase noise, we replace

$$e_A^2 = e_B^2 \rightarrow 4kT_{\text{bs}} \quad (\text{back-scatter equivalent temperature})$$

$$e_C^2 \rightarrow 4kT_{\text{eq}} \quad (\text{oscillator's equivalent temperature})$$

$$e_a^2 = e_b^2 = e_c^2 \rightarrow \frac{4}{3}kT_s \quad (\text{splitter, resistors are } R_0/3)$$

thus

$$S_{yx} = k \left(\frac{1}{4}T_{\text{eq}} - \frac{1}{4}T_s + T_{\text{bs}} \right) \quad (8.21)$$

Reminding that the available power is $P_0/4$ at each output, a naive choice for the estimator of the phase noise is

$$\widehat{S}_\varphi = \frac{\Re\{S_{yx}\}}{P_0} \quad (8.22)$$

¹⁰Keysight:E5500




Figure 8.7: (Replace $e_0 \rightarrow V_0$ for consistency to the clock signal, and also to avoid confusion between rms and peak values).

which gives

$$\widehat{S}_\varphi = \frac{k(T_{\text{eq}} - T_s + 4T_{\text{bs}})}{P_0} \quad (\text{affected by systematic error}) \quad (8.23)$$

This estimator is affected by a systematic error

$$\Delta S_\varphi = \frac{k(-T_s + 4T_{\text{bs}})}{P_0} \quad (\text{systematic error}) \quad (8.24)$$

with respect to the true $S_\varphi = kT_{\text{eq}}/P_0$. The systematic error is fixed by replacing the estimator with

$$\widehat{S}_\varphi = \frac{\Re\{S_{yx}\} + k(T_s - 4T_{\text{bs}})}{P_0} \quad (\text{estimator corrected for } T_s \text{ and } T_{\text{bs}}) \quad (8.25)$$

Unlike the directional coupler, there is no general rule to assess whether ΔS_φ is positive or negative.

8.2.4 V resistive power splitter (almost done)

Oscillator alone, terminated.

$$P_{\text{av}} = \frac{V_0^2}{8R_0} \quad (8.26)$$

$$S_{\text{rf}} = \frac{e_C^2}{4R_0} \quad (8.27)$$

$$S_\varphi = \frac{e_C^2}{V_0^2/2} \quad (8.28)$$

$$x = \frac{1}{4}e_C + \frac{3}{8}e_a - \frac{1}{8}e_b + \frac{5}{8}e_A + \frac{1}{8}e_B \quad (8.29)$$

$$y = \frac{1}{4}e_C - \frac{1}{8}e_a + \frac{3}{8}e_b + \frac{1}{8}e_A + \frac{5}{8}e_B. \quad (8.30)$$

$$xy = \frac{1}{16}e_C^2 - \frac{3}{64}e_a^2 - \frac{3}{64}e_b^2 + \frac{5}{64}e_A^2 + \frac{5}{64}e_B^2 \quad (8.31)$$

$$P_A = P_B = \frac{1}{32}V_0^2 = \frac{1}{4}P_{\text{av}} \quad (8.32)$$

Ideal case, no noise from the splitter ($e_a = e_b = 0$) and no back scatter ($e_A = e_B = 0$)

$$S_\varphi = \frac{e_C^2/16}{V_0^2/32} = \frac{e_C^2}{V_0^2/2} \quad (\text{as it should (not well presented)}) \quad (8.33)$$

In the real case

$$S_\varphi = \frac{\frac{e_C^2}{16R_0} - \frac{3e_a^2}{64R_0} - \frac{3e_b^2}{64R_0} + \frac{5e_A^2}{64R_0} + \frac{5e_B^2}{64R_0}}{\frac{V_0^2}{32R_0}} \quad (8.34)$$

Replace $e_a^2 = e_b^2 \rightarrow 4kT_s R_0$, where T_s is the splitter temperature. Replace $e_A^2 = e_B^2 \rightarrow 4kT_{\text{bs}} R_0$, where T_{bs} is the equivalent temperature associated to the back scatter, and we assume that the two detectors are equal. Thus

$$\frac{3}{64R_0} (e_a^2 + e_b^2) = \frac{3}{8} kT_s \quad (8.35)$$

$$\frac{5}{64R_0} (e_A^2 + e_B^2) = \frac{5}{8} kT_{\text{bs}} \quad (8.36)$$

$$\frac{V_0^2}{32R_0} = \frac{1}{4} P_{\text{av}} \quad (8.37)$$

and the error is

$$\Delta S_\varphi = \frac{-\frac{3}{8} kT_s + \frac{5}{8} kT_{\text{bs}}}{\frac{1}{4} P_{\text{av}}} \quad (8.38)$$

$$\Delta S_\varphi = \frac{k \left(\frac{3}{2} T_s + \frac{5}{2} T_{\text{bs}} \right)}{P_{\text{av}}} \quad (8.39)$$

Table 8.1: Phase noise analyzers.

Brand	Type	Technology	Splitter
Anapico	APPH series	Saturated mixer	3 dB coupler
Holme	(1)	Direct ADC	3 dB coupler
Holzworth	HA series	Saturated mixer	V resistive
Keysight	(all?)	Saturated mixer	Y resistive
Microchip	PhaseStation 53100	Direct ADC	Active (2)
NoiseXT / Arcale	DNA series	Direct ADC	3 dB coupler
NoiseXT / Arcale	NXA, DCNTS	Saturated mixer	3 dB coupler
Rohde Schwarz	FSWP, FSPN	SDR and down conv.	90° 3 dB coupler

(1) Andrew Holme, hobby project

(2) Jumpers on the front panel enable the use on an external splitter.

Open input

$$x = +\frac{1}{8}e_a - \frac{1}{4}e_b + \frac{3}{4}e_A + \frac{1}{4}e_B \quad (8.40)$$

$$y = -\frac{1}{8}e_a + \frac{1}{4}e_b + \frac{1}{4}e_A + \frac{3}{4}e_B \quad (8.41)$$

$$xy = -\frac{1}{16}(e_a^2 + e_b^2) + \frac{3}{16}(e_A^2 + e_B^2) \quad (8.42)$$

$$\frac{1}{16R_0}(e_a^2 + e_b^2) = \frac{1}{2}kT_s \quad (8.43)$$

$$\frac{3}{16R_0}(e_A^2 + e_B^2) = \frac{3}{2}kT_{bs} \quad (8.44)$$

and the error is

$$\Delta S_\varphi = \frac{-\frac{1}{2}kT_s + \frac{3}{2}kT_{bs}}{\frac{1}{4}P_{av}} \quad (8.45)$$

$$\Delta S_\varphi = \frac{k(-2T_s + 6T_{bs})}{P_{av}} \quad (8.46)$$

Table 8.1 summarizes the power splitters used in phase noise analyzers.

8.3 The weird case of the sub-thermal oscillator

8.4 Other effects

Crosstalk

Flicker noise

What else?

8.4.1 Correction accuracy and residual uncertainty

Temperature

Notice that $\Delta S_\varphi < 0$, which means that the instrument readout is always smaller than the DUT noise. This breaks the widespread belief that noise analyzers always over-estimate the DUT noise. Additionally, under-estimating the DUT noise may be embarrassing in some circumstances, like in laboratories of metrology.

Crosstalk

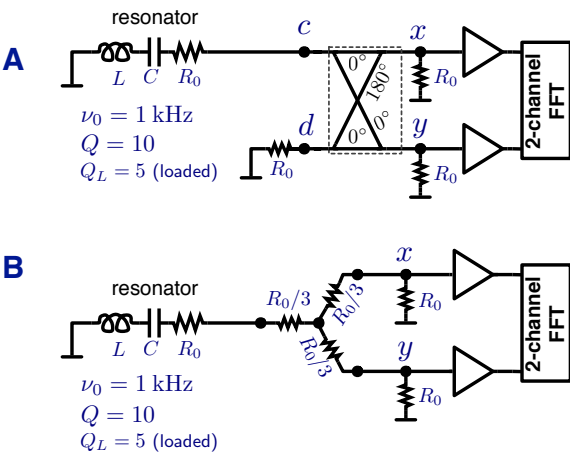


Figure 8.8: Experimental configurations A and B.




Figure 8.9




Figure 8.10




Figure 8.11




Figure 8.12




Figure 8.13




Figure 8.14




Figure 8.15




Figure 8.16




Figure 8.17




Figure 8.18

Chapter 9

Analog Phase Noise Analyzers

9.0.1 Saturated-Mixer Phase Noise Analyzer

Fig. 9.1 shows the block diagram of a traditional phase-noise analyzer. The instrument consists of two equal channels where a double balanced mixer compares the phase φ of the oscillator under test (DUT) to the phase θ and ψ of the references. The error signal sent to the LNA is equal to $k_d(\varphi - \theta)$, or to $k_d(\varphi - \psi)$, where k_d is the mixer's phase-to-voltage gain. The typical value of k_d is between 0.1 V/rad and 1 V/rad, depending on the signal level, technology, and frequency.

This operating mode requires that $\varphi - \theta = \pi/2$ within ± 0.1 rad, and likewise $\varphi - \psi$, which is ensured by the PLL. The error signal, corrected for the equation of the PLL, is proportional to the DUT phase fluctuation vs the references. A problem with mixers is that the power range is rather narrow, from a minimum of ≈ 5 dBm to a maximum of 17–20 dBm, depending on technology. The power range can be extended adding an amplifier at each RF input.

Assuming that the two channels are independent (separate mixers and separate synthesizers, and no crosstalk), the average cross spectrum is proportional to the DUT phase noise because the background noise of the individual channel is rejected by the averaging process. It is worth mentioning that averaging is essential to reject the noise of the synthesizers. The synthesizers are necessary for the instrument to be suitable to a wide range of frequency, but they are generally noisier than oscillators.

Instruments of this type are commercially available from Anapico, Berkeley Nucleonics Corp., Holzworth, Keysight Technologies, NoiseXT/Arcale, and Quantic Wenzel (a brand of Wenzel Associates). The frequency synthesizers may be inside the

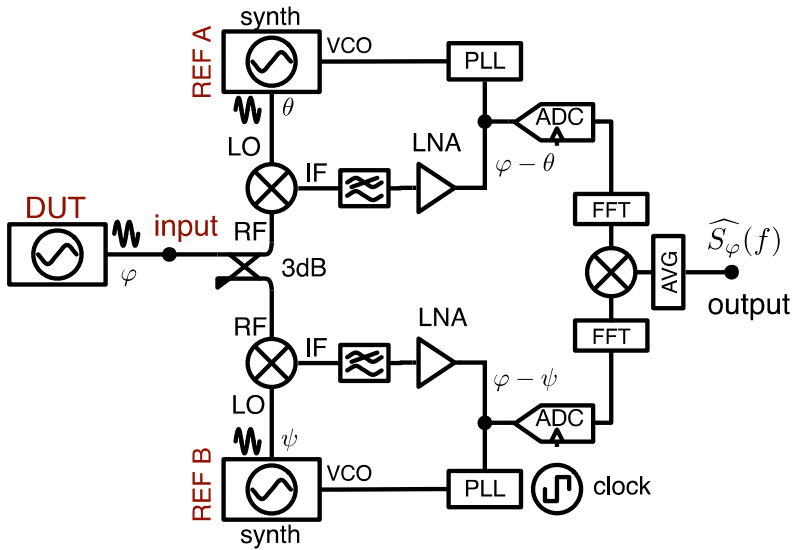


Figure 9.1: Traditional phase noise analyzer. The mixers, saturated at both input, work as phase-to-voltage converters. Two equal channels are used to reject the background noise after averaging on multiple cross spectra.

instrument, or external modules. The input power splitter may differ from the 3 dB directional coupler shown in Fig. 9.1. We have seen Y resistive networks (Keysight), and V resistive network (Holzworth).

Chapter 10

Digital Phase Noise Analyzers

ZZZZZZZZZZZZ

Chapter 11

Analyzers

11.1 **Aeroflex**

The old PN9000

11.2 **Anapico**

11.3 **Anritsu**

11.4 **BNC**

11.5 **Andrew Holme's hobby project**

Copied almost verbatim from <http://aholme.co.uk>, and from various email exchanges

11.5.1 Architecture

This apparatus measures phase noise down to wideband noise floor levels below -170 dBc/Hz. Historically, such measurements were either difficult or expensive to make. Based on the work [1] of Grove et al., the method described here is cheap, sensitive, accurate and requires no calibration. It is a differential measurement between a device under test (DUT) and a reference oscillator, which are connected to the SMA ports at the extreme left. Ideally, the reference oscillator should be an order




Figure 11.1

of magnitude (or more) quieter, so DUT noise predominates.

The largest board is a Xilinx SP605 FPGA evaluation kit, attached to which (by the FMC connector) is a Linear Technology DC1525A-A quad ADC evaluation board¹, minus its top-right corner. The two identical boards to which the DUT and reference are connected are custom-made power splitters, feeding four ADC inputs through coaxial cables. The part-populated yellow board, spare from my GPS project, was re-purposed as a clean power supply for a Crystek CPRO33-77.760 crystal oscillator, which can just be seen mounted vertically on the ADC clock input SMA.

ADC resolution is 14-bits. Samples are received by the FPGA on eight LVDS pairs, each of which can run at up to 1 Gbps. In the FPGA, the four ADC channels are digitally down-converted to IQ complex baseband, low pass filtered, decimated, and phase-demodulated using a CORDIC arc-tangent block. The resultant phase data is streamed to a Windows laptop via UDP over Ethernet. Phase data is converted to power spectral density by Fast Fourier Transform (FFT) and displayed in real-time. The quality of the measurement is refined by correlation / averaging the longer it runs. The inputs are down-converted by mixing with a quadrature NCO. NCO rate is set close to input frequency, but cannot remain exactly equal, so the baseband signal is close to but rarely at zero frequency. The CORDIC output is phase noise superimposed on a gentle linear ramp, the gradient of which depends on the small unavoidable difference between input frequency and NCO rate.

11.5.2 Cancel and correlate

It's possible to measure phase noise of a DUT, relative to the sampling clock, using a single ADC channel. In fact, I did so as a first step. So why the two inputs and four ADC channels? Because using a single 14-bit ADC, even assuming zero sampling clock

¹Andrew informed me (May 29, 2023) that the DC1525A-A evaluation board is no longer available. Albeit the Gerber files can be downloaded from the AD web site, the chips themselves are currently out of stock. The demo board is based on the LTC2175-14 chip, a quad 14-Bit ADC sampling at 125 Msps, 5-170 MHz input, 1.8 V supply.

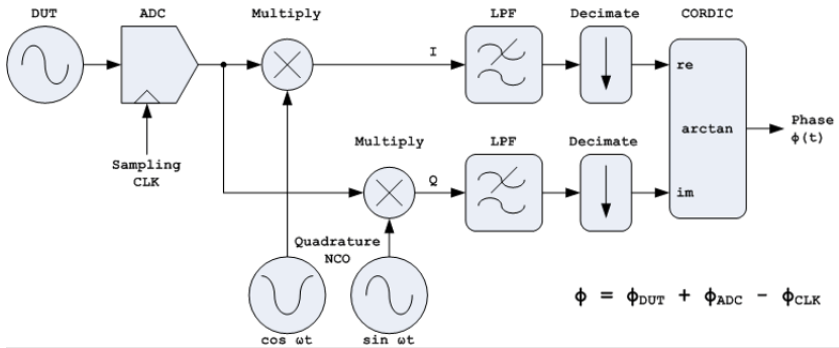


Figure 11.2

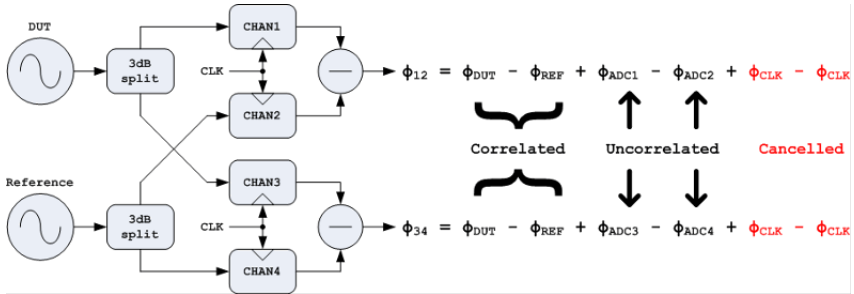


Figure 11.3

jitter, the simulation-predicted noise floor is worse than -150 dBc/Hz. Typically, more like -140 dBc/Hz. Grove et al. did two things to get the noise floor down below -170 dBc/Hz.

The first neat trick is to make a differential measurement between DUT and reference, by subtracting their phases. The phase noise of the sampling clock is cancelled-out. The quality of the measurement depends on the reference, not the ADC clock, which can be slightly inferior. Although phases are subtracted, uncorrelated noise powers add. The result is the sum of DUT and reference noise. The reference can be an order of magnitude quieter than the DUT, or, if two identical oscillators are compared, half the power (-3 dB) can be attributed to each.

Cancelling clock noise still leaves ADC quantization noise and thermal noise, which raise the noise floor above -150 dBc/Hz, as already mentioned. The second neat trick is to duplicate the measurement. This is the reason for the power splitters. Everything in the FPGA (and downstream in software) is replicated. We simultaneously make two completely independent measurements. Both contain the same (correlated) DUT+Reference noise; but different (uncorrelated) ADC noises. The latter are then greatly attenuated using the cross spectrum experimental method described in [2] by E. Rubiola and F. Vernotte.

DUT and reference do not have to be the same frequency. If the sampling points




Figure 11.4

of all four ADC channels are displaced by a small time offset Δt due to jitter on the common sampling clock, it affects the measured phases in proportion to their frequencies: $\Delta\theta = \omega\Delta t$. Reference phase must be scaled by the DUT/REF frequency ratio to cancel clock jitter. ADC noises remain uncorrelated, scaled or not. It can be beneficial to use a higher reference frequency, because its phase noise contribution, Φ_{REF} , will then be scaled down. All we want from it is stability.

11.5.3 Phase wraps

The 32-bit fixed-point binary output of the CORDIC is encoded in semi-circles with 31 bits after the binary point. The most significant bit represents ± 1 semi-circle or $\pm\pi$ radians. The angle can be interpreted as signed or unsigned, without ambiguity. Due to unavoidable frequency differences, the measured phases will always be slowly but steadily ramping, either up or down. Phase ‘wraps’ occur every 2π radians. Real measurements have phase noise superimposed on the ramp and can “chatter” back and forth several times as they go through the wrapping point.

These discontinuities must be removed before the FFT. One way is to unwrap them. Integer precision can be extended before the binary point and 1.0 circles (2.0 semi-circles) are added or subtracted every time the raw data wraps. Another way — the method I am using — is to convert phase to instantaneous frequency, ω , by differentiating. Each phase wrap generates a carry or a borrow from the differentiator; but there are no jumps or glitches in its 32-bit output. The FFT then calculates power spectral density (PSD) of frequency fluctuation, which is transformed to PSD of phase afterwards.

An ideal differentiator converts $\sin(\omega t)$ to $\omega \cos(\omega t)$ and so has gain ω . The 90 degree phase shift doesn’t affect power. Originally, LivePlot divided output bins by ω^2 to obtain PSD of phase; however, this caused a slight droop in wideband noise floors approaching Nyquist, where gain of the first-difference digital differentiator deviates from ideal. This small inaccuracy was corrected in LivePlot version 16 by using actual gain of the digital differentiator instead of ω .

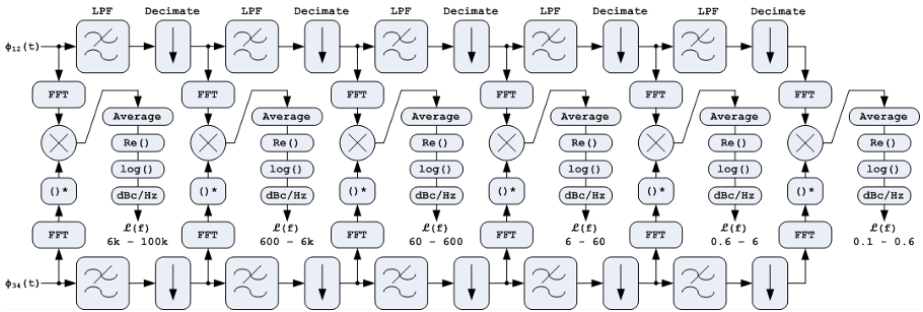


Figure 11.5

11.5.4 Frequency domain

The final output, $\mathcal{L}(f)$, is the “one-sided” power spectral density of phase, relative to the carrier, in dBc/Hz. This is plotted on a log/log scale, typically over the frequency range 0.1 Hz to 100 kHz. Samples are transformed into the frequency domain by Fast Fourier Transform. Making two independent measurements doubles the amount of processing required. In order to execute this in real-time, with the same number of data points in each decade, the processing is broken up into a succession of identical stages, decimating by 10 at each step:

The above diagram includes cross-correlation. Complex FFT outputs bins from one measurement are multiplied by the complex conjugate of the corresponding bin from the other. The complex product is averaged for minutes, hours or days. Uncorrelated noises are attenuated by $5\log_{10}(N)$ where N is the number averaged. After a while, the average real part tends to contain only correlated noise. The average imaginary part is an indicator of system noise floor. 5dB steps can be seen in the imaginary part between decades because of decimation. Each decade averages 10 times more data than the next lower in frequency.

ADC sampling rate is 77.76 Msp/s, decimated to 607.5 kps in the FPGA. FFT length is typically 1000 points, and bins 10 to 99 are plotted from each decade. Length can be increased to get finer detail. Bins $LEN/100$ to $LEN/10$ are always plotted in the middle decades; but more are required (up to $LEN * 100/607.5$) in the first stage to reach 100 kHz. Fewer are needed in the last stage because the graph starts from 0.1 Hz. Low pass filters are 6th order Butterworth IIR with a normalised cut-off frequency of 20/1000.

Results must be scaled to compensate for CIC pass-band droop; differentiator and LPF gain; FFT growth; FFT bin size (converting to dB/Hz) and equivalent noise bandwidth (leakage) of a raised cosine window function, which is applied before the FFT. For efficiency, combined adjustments are pre-computed in dB and applied as offsets to the final outputs.

11.5.5 Spurs

Does this sound too good to be true? Are you wondering what the snag is? The biggest problem is instrument-generated spurs, due to ADC quantization. These appear when harmonics of the input signal fall close to harmonics of the sampling clock. I discovered them the hard way. The DC1525A ADC board works up to 125 Msps and I was making measurements of 5 MHz sources using a 124.998 MHz sampling clock! None of the papers I read beforehand led me to expect this, and I was puzzled for a while, until I reproduced the spurs using a very simple simulation. Moving to 77.76 Msps improved matters.

The same signal can be fed to both the DUT and reference inputs with a third power splitter. I used this configuration to estimate system noise floor, before I knew about the imaginary part of the averaged cross-product. When I started making differential measurements, I discovered another problem: low-frequency spurs, due to crosstalk, at the DUT – Reference difference frequency. Fortunately, although they are quite noticeable in the imaginary part of the cross-product, these spurs barely push through into the real part.

11.5.6 Examples of results

Each of the plots below was averaged over several hours. DUT and reference inputs were 5 MHz. Blue trace is $\mathcal{L}(f)$, the one-sided power spectral density of phase, as estimated by the real part of the average cross-product. Green trace is system noise floor in the imaginary part. The Wenzel ULN is an ultra-low-noise reference oscillator:

The Wenzel 500-03220 is about 1 Hz off-frequency, judging by the low-frequency crosstalk spurs around 1 and 2 Hz, evident most strongly in the imaginary part. All plots have spurs at the 50 Hz power line frequency. Harmonics of 50 Hz and other "real" spurs are visible on the synthesizer plots. The -160 dBc spur around 21 kHz in (a) is $77759000 \times 58 - 5000001 \times 902$. The Dana 7020 "Digiphase" synthesizer's closed loop response matches figure 26-8 in Garry Gillette's essay on page 290 of "Analog Circuit Design" edited by Jim Williams.

My 2019A takes time to warm-up before delivering this performance. Notice how it selects a narrower loop bandwidth when modulated. Peak phase deviation of 0.001 radians is -66 dBc. Spur amplitudes cannot be read directly off the dBc/Hz scale. The plot is scaled by FFT bin size to normalized power spectral density in a 1 Hz bandwidth. Clicking on a spur in the GUI reveals its true amplitude. LivePlot reports 5.81 kHz and -66.3 dBc. Break through from the audio oscillator is observable with modulation disabled.

11.5.7 Acknowledgments

I wish to thank Joël Imbaud of Institut FEMTO-ST who replicated the SP605/DC1525A setup and verified its accuracy against a commercial Symmetricom 5120A. Instead of custom boards, Joël used Mini-Circuits attenuators; ZMSC-2-1+ connectorized power splitters; SPB10.7+ and SLP5+ filters; and AMP-77 amplifiers. He removed the SFP cage to avoid sawing his ADC board. I checked the Gerbers before cutting mine.




Figure 11.6

I also wish to thank Adrian Rus who obtained good agreement between his SP605 / DC1525A and a Rohde and Schwarz FSWP for one Wenzel ULNO; and against the manufacturer's specification of another. Adrian measured DUT noise of -176 dBc/Hz with a system noise floor below -180 dBc/Hz after cross correlation. He tested many ways of feeding the ADC and recommends this configuration for minimizing artifacts such as negative bias, and reaching ultra-low noise floors, even when measuring low output level DUTs:

Two high-quality reference oscillators are connected directly (via attenuators) to the ADC inputs. Their noises are removed by cross-correlation and the measurement is DUT noise only. Adrian modified the 1525A clock input to accept sinusoidal drive and reduced clock noise by soldering $50\ \Omega$ resistors across the ADC inputs. He found 6dB resistive splitters superior to 3dB reactive types. Clean, independent power supplies are mandatory for the amplifiers. Adrian Rus is co-author of a paper [3] entitled “Artifacts and Errors in Cross-Spectrum Phase Noise Measurements” concerning how the thermal energy of the power splitter has opposite phases on the two DUT channels, producing an anti-correlated bias, leading to artificially low measurements of b_0

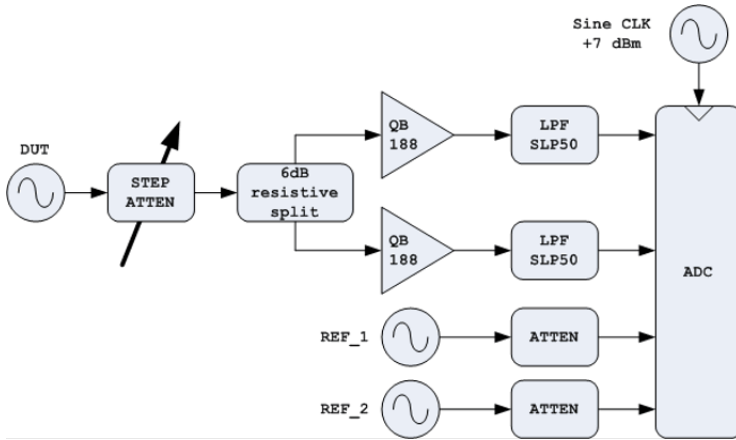


Figure 11.7




Figure 11.8

wideband noise.

11.5.8 Windowing and overlapping

Andrew uses raised cosine window function. The figure shows a 6 second measurement of a -170 dBc/Hz signal, no overlap (top), and 50% overlap (bottom). Such overlap gives all samples exactly equal weighting.

11.5.9 Multiplier experiment (from Andrew's email)

ADC clock rate is 132.9277 MHz and hardware decimation is 64. There is software decimation by 10 between each FFT decade, as you know.

Samples per second at FFT decade which covers 1 Hz = $132.9277\text{e}6/64/1\text{e}5 = 20.769953$

FFT segment end points are

$132.9277\text{e}6/64/1\text{e}5 * 5/128 = 0.8113263$

$$132.9277\text{e}6/64/1\text{e}5 * 50/128 = 8.1132629$$

Despite poor resolution, I must use FFT length 128 to get as many averages as possible.

$$\text{Averages per second} = 132.9277\text{e}6/64/1\text{e}5/128 = 0.1622653$$

$$\text{Runtime for } 1\text{e}4 \text{ averages} = 1\text{e}4/0.1622653/60/60 = 17 \text{ hours}$$

11.6 Holzworth

11.6.1 Contacts at Holzworth

- Jason Breibarth left mid 2020. Private address Jason.breibarth@gmail.com
- Nicolas Laves, nlaves@wtcom.com, application engineer, Wireless Telecom Group wireless telecom group.com P: +1 (303) 325-3473 ext 117, Boonton, Com-Magility, Holzworth, Microlab, Noisecom
- Matthew Diessner, mdiessner@wtcom.com

11.6.2 Emailing with Jason Breibarth, January 2019

The input splitter is resistive. It's a two resistor, 50 Ω each leg (like used on power detector/ALC loops). The mixer RF port is not saturated, except at high input power levels. We have an input power range of 0-20 dBm. 10-12 dBm is the optimal point for linear operation of the mixer and highest dynamic range. The LO port is driven at normal drive operation (saturated). Both the LO and RF input points are buffered with a very low phase noise amplifier so keep any reflections and to maintain isolation between channels. One of the problems with the setups I've seen is the Wilkinson goes straight into the mixer RF ports and those ports are typically terrible return loss — as bad as -3 dB in cases. The isolation of the Wilkinson goes down to 3 dB in most cases.

11.6.3 HA7402C

11.6.4 HA7062C

11.6.5 HA7062D

11.6.6 Measurement of a two-port DUT

11.6.7

11.7 John Miles and Jackson Labs

11.7.1 John Miles wrote, July 17, 2021

I was hoping to be able to use dual ADCs, and in fact have built some other gear with AD9650s, but for the 53100A I ultimately decided to stick with single-channel parts (AD9265). The crosstalk issue was the #1 reason. The dual-channel parts work very well, but there is just no way to achieve the isolation needed for clean ADEV plots when the sources are separated by a few Hz.




Figure 11.9

Another issue is that the dual-channel parts have their own internal clock buffers. As a result, a four-channel instrument built with a pair of two-channel ADCs won't really be able to offer four independent channels, since you end up measuring the (substantial) noise of that clock buffer if you configure a cross-spectral measurement that relies on averaging between ADCs in separate chips (i.e., if channels 0 and 1 correspond to one physical ADC package and channels 2-3 correspond to the other, then 0-1 x 2-3 works well, but 0-2 x 1-3 is not useful.)

11.8 Keysight E5052B, optional E5053A

11.9 Keysight, the new analyzer

11.10 Koheron 250 MSPS acquisition board

<https://www.koheron.com/fpga/alpha250-signal-acquisition-generation>

The ALPHA250 is a programmable board built around a Zynq 7020 SoC. It features a 100 MHz RF front end with 14-bit ADCs and 16-bit DACs, at 250 MSPS. The RF channels are clocked by a dual PLL, ultra-low jitter clock generator. It includes a 4-channel 24-bit ADC and a 4-channel 16-bit DAC. The board comes with a comprehensive, open source, FPGA / Linux reference design.

Koheron has a phase noise analyzer project on GitHub at the URL <https://github.com/Koheron/koheron-sdk/tree/master/examples/alpha250/>

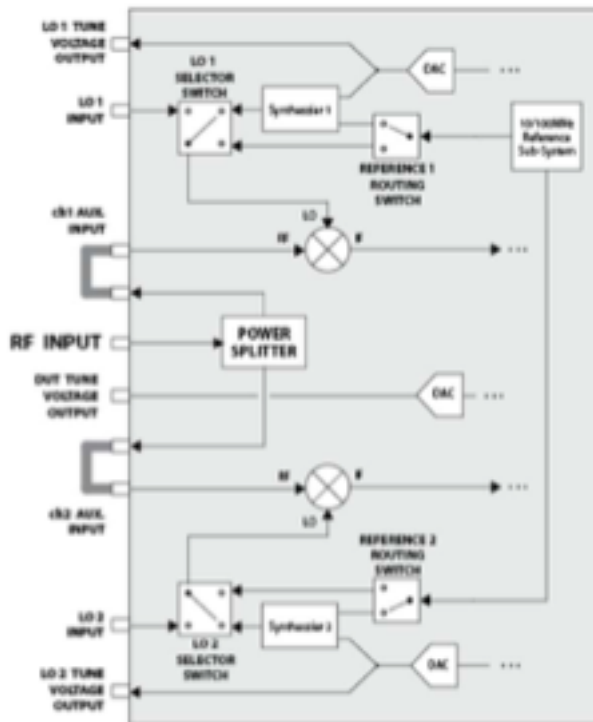
HA7062C (front end)

Figure 11.10

phase-noise-analyzer.

11.11 Meir Alon's project

Just started, August 2023, meir.alon@weizmann.ac.il Meir attended the EFTS

11.12 Microchip 5120A

11.13 Microchip 5125A

11.14 John Miles' TimePod

11.15 Rohde Schwarz FSWP and FSPN

The reference article is Feldhaus G, Roth A - A 1 MHz to 50 GHz Direct Down-Conversion Phase Noise Analyzer with Cross-Correlation - Proc EFTF 2016.

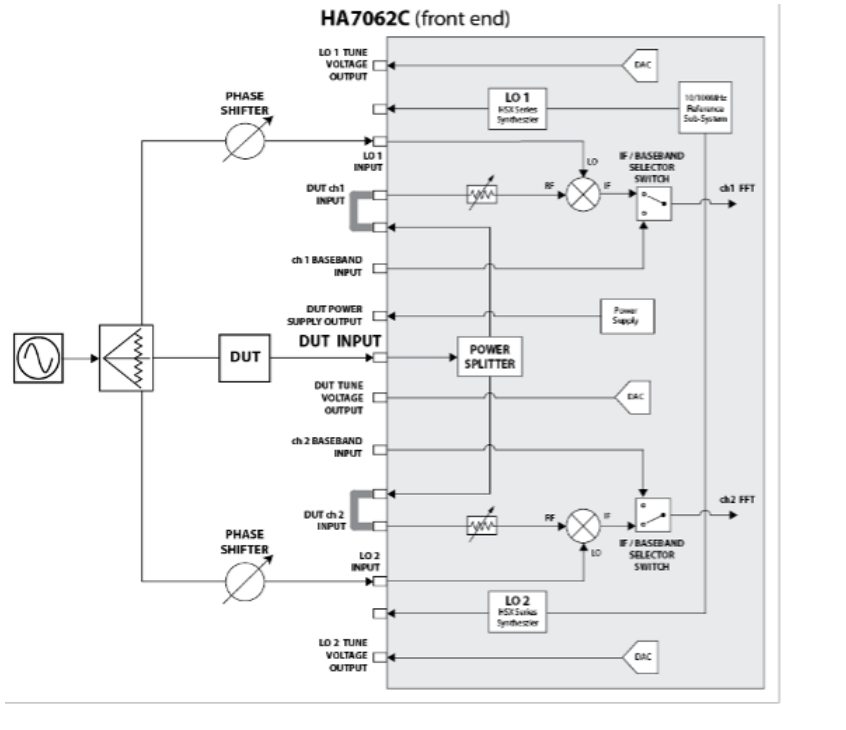


Figure 11.11

11.15.1 Alexander Roth, January 2019

In the FSWP we use 3 dB reactive power splitters to split the input signal. Three different splitters are used for the frequency ranges 1MHz to approximately 1 GHz, 1 GHz to 8 GHz and 8 GHz to 50 GHz. The isolation between the two receiving paths is therefore better as with a resistive divider but of course not perfect. I expect that we are in the range of -20 dB of decoupling between the two outputs of the splitters. But this is only true if the DUT has good output matching, what's not very likely for oscillators without buffer amplifier. Therefore, an attenuator at the RF input of the FSWP can improve the decoupling between the two channels and therefore also lower coupling from noise from one channel to the other.

The cross spectrum is estimated with $S_{yx} = |YX^*|$

The FSWP does not use a zero IF concept to measure PN and additionally the IF's of the two receiving path are slightly different (< 1 Hz offset). Therefore, there is no fix phase relation between the two channels.

We use two different IF modes to measure PN:

Frequency Offset < 1 MHz \implies IF = 1.3 MHz

Frequency Offset > 1 MHz \implies IF = < 10 Hz

The IQ outputs of the analog mixer are sampled with 18 Bit ADCs and converted to zero IF in the digital domain within a FPGA.

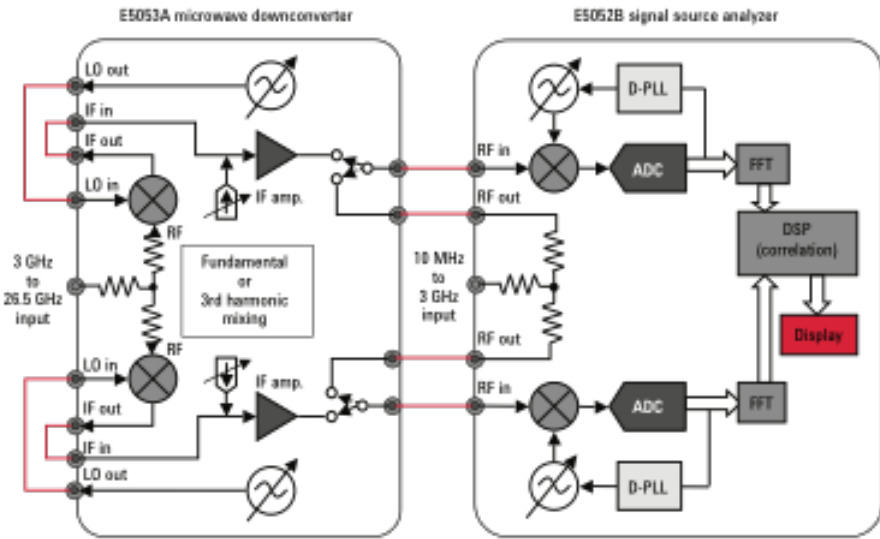


Figure 11.12




Figure 11.13: Noise floor for the 4 ADC version (ALPHA250-4) inputs (50 Ω terminated), from https://assets.koheron.com/images/products/alpha250/ALPHA250-4_noise-floor.png?03052024090915.

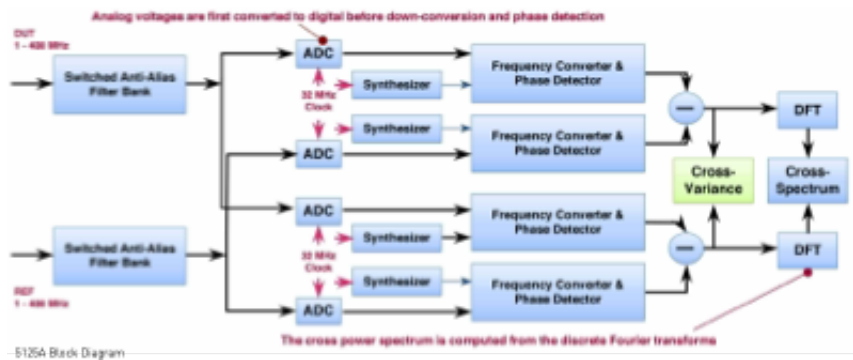


Figure 11.14

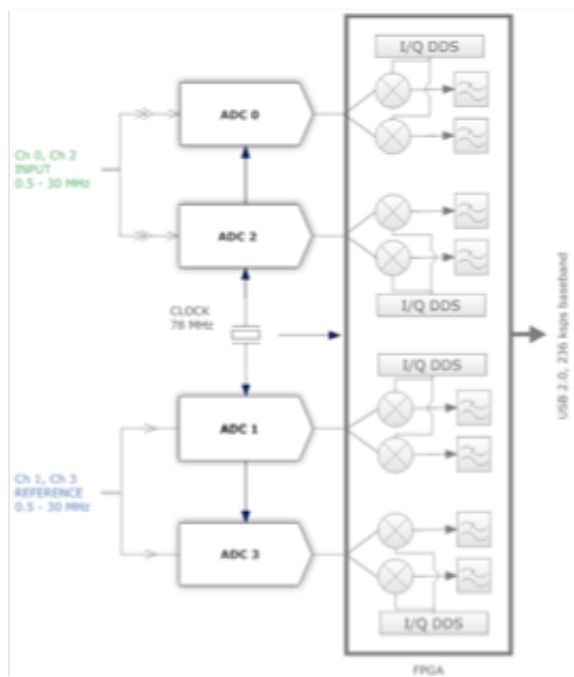


Figure 11.15

Phase- and/or Amplitude-Demodulation is all digital.

11.15.2 FSPN

Basically, a cheaper version of the FSWP, where numerous functions that are not strictly related to phase noise have been removed.

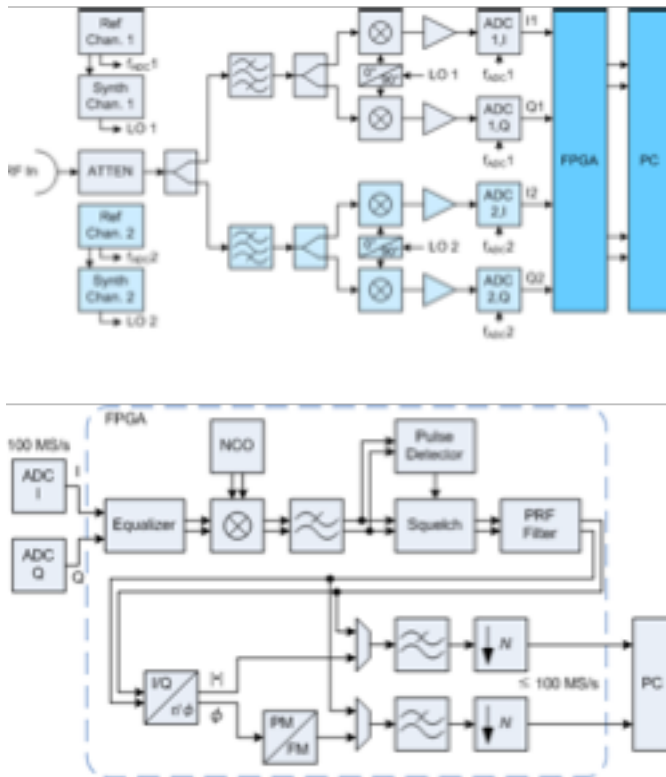


Figure 11.16: Block diagram of the FSWP (top), and the detail of the FPGA (bottom). Reproduced from Feldhaus G, Roth A - A 1 MHz to 50 GHz Direct Down-Conversion Phase Noise Analyzer with Cross-Correlation - Proc EFTF 2016.

11.16 NoiseXT (Sphera Arcale)

11.16.1 Contacts

ADRIEN Laurent laurent.adrien@sphera.com

Bruno Cohen b.cohen@arcale.net

Sylvain DAVIET s.daviet@arcale.net

11.16.2 DCNTS

11.16.3 DNA400M

Disassembling a specimen, we saw a TI ADS42LB69

Discussion with Laurent Adrien (confidential)

1. Possibilité de spectre Autocorrélation sur les canaux A et B. OK

- Implémenté. vous avez la possibilité de sélectionner, soit Xcorr (inter-

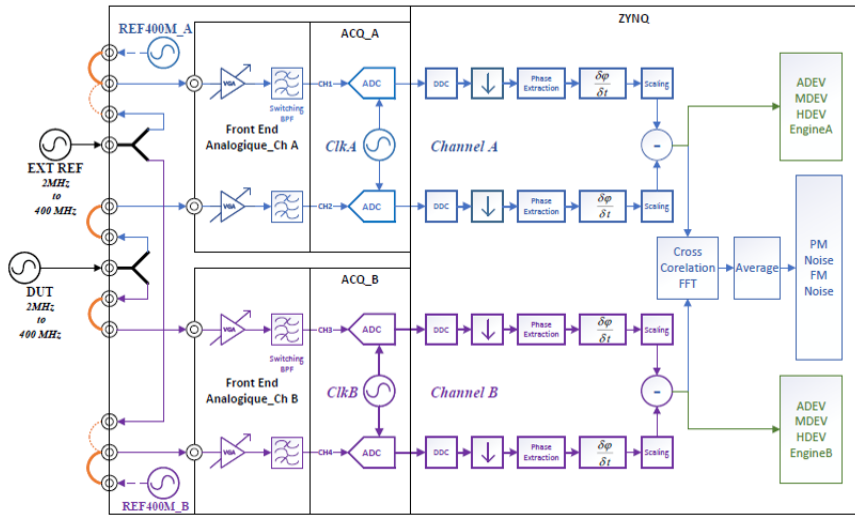


Figure 11.17: DNA400M. **Laurent Adrien, personal communication, requires permission.**

corrélation A&B) ou Autocorr A ou B (autocorrélation par canal)

2. Raw Data OK

- La cadence des RAW data est déterminé par le ENBW sélectionné
- 6 fichiers :

☒ DUT1.txt
☐ DUT2.txt
☐ fractional_freq_path1.dat
☐ fractional_freq_path2.dat
☐ REF1.txt
☐ REF2.txt

- 4 fichiers unwrapped phase, un fichier par voie (DUT1, DUT2, REF1, REF2)
- 2 fichiers fractional frequency , un fichier par canal A et B. (fractional_freq_path1 & 2)
- Les fichiers fractional-freq_path1&2 seront renommés en A&B, petit bug
- Les fichiers fractional-freq_path sont compatible stable32

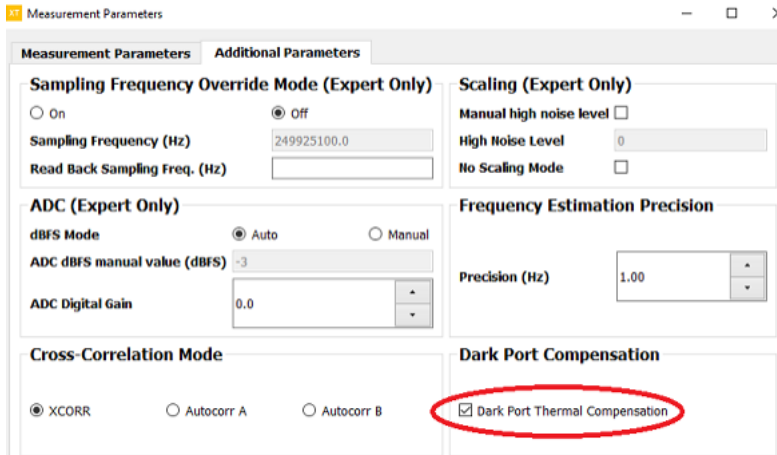
3. Possibilité d'avoir 4 fréquences différentes sur les 4 voies. En cours

- Dans le principe cela est possible. les fréquences sont calculé en entrée ADC sur chaque voie.

- pour les voies références pas de problème, cela fonctionne.
- cela doit être testé et validé pour les voies DUT 1 & 2 .

4. Power splitter dark port compensation. OK

- Compensation du bruit dark port du power splitter de la voie DUT.



5. Inter-corrélation Phase noise identification partie réelle négative. A voir

- Pas encore fait

6. Phase noise, Down load trace data $|Real|$ et $|Img|$: OK

- possible de sauvegarder les résultats de mesure de bruit de phase au format csv.
 - Menu Utility/Trace/Save trace
- Possible de sauvegarder Noise spectrum
 - Menu Utility/Trace/Trace select “noise spectrum”
 - Si l’estimateur utilisé est $|R|$ alors vous aurez $|R|$ en Log
- Possible de sauvegarder $|img|$
 - Menu Utility/Trace/Trace select “system residual”
 - Par défaut le system residual est smoothed, si vous le souhaitez, vous pouvez choisir un system residual no smoothed.

11.17 tinyPFA

<https://www.tinydevices.org/wiki/pmwiki.php?n=TinyPFA.Homepage>

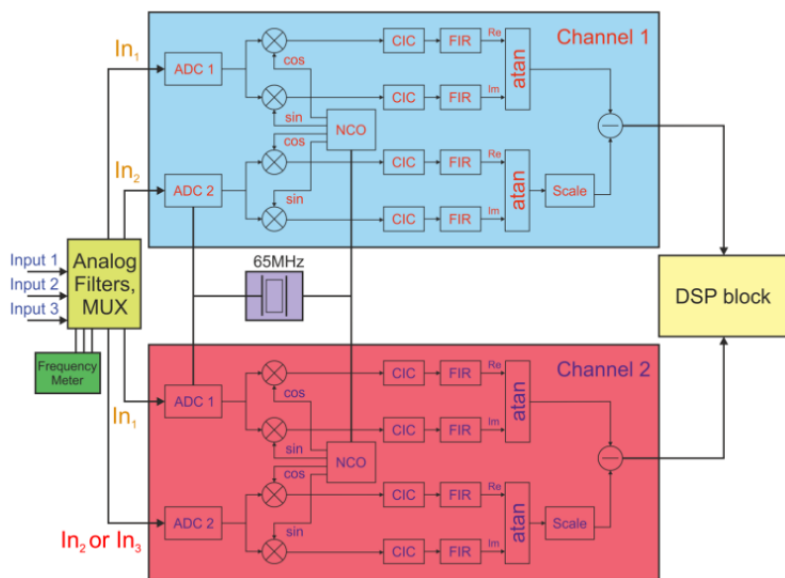


Figure 11.18: Fully digital, 1-100 MHz carrier frequency, 100 kHz max Fourier frequency. Not sure that this is from Vremya.

11.18 Vremya VCH-314/315/315M/325

Regardless of the difficulty to get information, I may decide to boycott Vremya

<https://www.vremya-ch.com/index.php/en/category/products-en/freq-comparators-en/index.html>

11.19 Yanjun Ma BG6KHC's project

All this comes from <https://www.qsl.net/bg6khc/>, email pn2060a@gmail.com and pn2060a.taobao.com

11.19.1 PN2060A/B Phase Noise Analyzer

https://qsl.net/bg6khc/pn2060ab_phase_noise_analyzer.htm

11.19.2 PN2060C Phase Noise Analyzer

https://qsl.net/bg6khc/pn2060c_phase_noise_analyzer.htm

11.19.3 Andrew “Drew” Wollin VK4ZXI's comments

<https://vk4zxi.blogspot.com/2023/07/an-economical-way-to-measure-phase.html>, thanks to Jean-Michel Friedt for informing me.

The PN2060A Phase Noise Analyzer is a low-cost instrument from China, currently USD 480. It is available from its maker, Yanjun Ma BG6KHC. Yanjun responds to emails very quickly and is very helpful.

It works in a similar way to the Microsemi 53100 phase noise analyser (USD 24000), itself a modern lower-cost instrument compared to those from Keysight, R&S and others.

PN2060A Phase Noise Analyzer <https://qsl.net/b/bg6khc/>. The location of the instructions and programs <https://qsl.net/b/bg6khc//pn2060a/>. The Microsemi 53100 phase noise analyser, <https://www.microsemi.com/product-directory/phase-noise-and-allan-deviation-testers/5565-53100a#overview>. Very good application notes on using a phase noise analyser <https://www.microsemi.com/product-directory/phase-noise-and-allan-deviation-testers/5565-53100a#resources>.

The analyser is a work in progress by the maker. Currently, there are two models, PN2060A and a later PN2060B. I have a PN2060A. I am not sure of the availability or pricing for the PN2060B, other than more expensive. The PN2060B has a lower noise floor compared to the PN2060A. There is a PN2060C in development, according to the maker's website. The PN2060C “Eliminate the ‘phase compensation,’ and minimize the coupling among ADC channels.” The PN2060A Phase Noise Analyzer, its architecture and block diagram. Two external reference oscillators and a split signal from the DUT go into separate ADCs, with the mixing and filtering done digitally before it all passes by a USB3 link to a computer for final processing.

Chapter 12

Photonic Systems¹

Applications of the fiber²

Science and technology require reference microwave sources with ever increasing stability and spectral purity, and of course suitable measurement systems. The frequency synthesis from the HF/VHF region is no longer satisfactory because of the insufficient spectral purity inherent in the multiplication, and because of more trivial limitations like mass, volume and complexity. The domain of RF/microwave photonics is growing fast [?]. The generation of microwaves from optics, or more generally ‘with optics,’ is providing new solutions. Examples are mode-locked lasers [?, Jones et al., 2003], opto-electronic oscillators photonic [?, ?] (OEO), optical micro-cavities [?], optical parametric oscillators [?, ?] (OPO), and frequency combs [?].

The optical fiber used as a delay line enables the generation [?, ?] and the measurement [?] of stable and highly spectrally pure microwave signals. The optical fiber is a good choice for the following reasons.

1. A long delay can be achieved, of 100 μ s and more, thanks to the low loss (0.2 dB/km at 1.55 μ m and 0.35 dB/km at 1.31 μ m).
2. The frequency range is wide, at least of 40 GHz, still limited by the opto-electronic components

²Most of this chapter comes from K. Volyanski et al., JOSAB. . .”

3. The background noise is low, close to the limit imposed by the shot noise and by the thermal noise at the detector output.
4. The thermal sensitivity of the delay ($6.85 \times 10^{-6}/\text{K}$) is a factor of 10 lower than the sapphire dielectric cavity at room temperature. This resonator is considered best ultra-stable microwave reference.
5. In oscillators and phase-noise measurements, the microwave frequency is the inverse of the delay. This means that the oscillator or the instrument can be tuned in step of 10^{-5} – 10^{-6} of the carrier frequency without degrading stability and spectral purity with frequency synthesis. Finer tuning is possible at a minimum cost in terms of stability and spectral purity.

We describe a series of experiments related to the generation and to the measurement of low-phase-noise microwave signals using a delay implemented with an intensity-modulated beam propagating through an optical fiber. A 10 GHz oscillator prototype exhibits a frequency flicker of 3.7×10^{-12} (Allan deviation) and a phase noise lower than -140 dB rad^2/Hz at 10 kHz off the carrier. The same values are achieved as the sensitivity in phase noise measurements in real time. This sensitivity is sufficient to measure the frequency flicker of a sapphire oscillator, which is considered the reference in the field of low-noise microwave oscillators [?]. In phase noise measurements, enhanced sensitivity is obtained with the cross-spectrum method, which takes correlation and averaging on two independent delay lines.

This article stands on Ref. [?] for the oscillator, and on Ref. [?, ?] for the measurements. We aim at providing practical knowledge, adding engineering, accurate calibration, the analysis of the cross-spectrum technique, the phase-noise model of the oscillator, and experimental tests.

12.1 Basic concepts

12.1.1 Phase noise

Phase noise is a well established subject, clearly explained in classical references, among which we prefer [?, ?, ?, ?] and [?, vol. 1, chap. 2]. The quasi-perfect sinusoidal signal of frequency ν_0 , of random amplitude fluctuation $\alpha(t)$, and of random phase fluctuation $\varphi(t)$ is

$$v(t) = [1 + \alpha(t)] \cos [2\pi\nu_0 t + \varphi(t)] . \quad (12.1)$$

We may need that $|\alpha(t)| \ll 1$ and $|\varphi(t)| \ll 1$ during the measurement. The phase noise is generally measured as the average PSD (power spectral density)

$$S_\varphi(f) = \langle |\Phi(jf)|^2 \rangle_m \quad (\text{avg, } m \text{ spectra}), \quad (12.2)$$

The uppercase denotes the Fourier transform, so $\varphi(t) \leftrightarrow \Phi(jf)$ form a transform inverse-transform pair. In experimental science, the single-sided PSD is preferred to the two-sided PSD because the negative frequencies are redundant. It has been

found that the power-law model describes accurately the phase noise of oscillator and components

$$S_\varphi(f) = \sum_{n=-4}^0 b_n f^n \quad (\text{power law}) \quad (12.3)$$

coefficient	noise type
b_{-4}	frequency random walk
b_{-3}	flicker of frequency
b_{-2}	white frequency noise, or phase random walk
b_{-1}	flicker of phase
b_0	white phase noise

The power law relies on the fact that white (f^0) and flicker ($1/f$) noises exist per-se, and that a phase integration is present in oscillators, which multiplies the spectrum $\times 1/f^2$. If needed, the model can be extended to $n < -4$. Of course, the power law can also be used to describe the spectrum the fractional frequency fluctuation $y(t) = \dot{\varphi}(t)/2\pi\nu_0$

$$S_y(f) = \frac{f^2}{\nu_0^2} S_\varphi(f) = \sum_{n=-2}^2 h_n f^n. \quad (12.4)$$

Another tool often used is the two-sample (Allan) variance $\sigma_y^2(\tau)$ as a function of the measurement time τ . Notice that the symbol τ is commonly used for the measurement time and for delay of the line (Sec. 12.1.3). We will added a subscript when needed. For the most useful frequency-noise processes, the relation between $\sigma_y^2(\tau)$ and $S_y(f)$ is

$$\sigma_y^2(\tau) = \begin{cases} \frac{h_0}{2\tau} & \text{white frequency noise} \\ h_{-1} 2 \ln(2) & \text{flicker of frequency} \\ h_{-2} \frac{(2\pi)^2}{6} \tau & \text{frequency random walk} \end{cases} \quad (12.5)$$

12.1.2 Cross-spectrum method

Inevitably, the measured noise is the sum of the device-under-test (DUT) noise and of the the instrument background. Improved sensitivity is obtained with a correlation instrument, in which two separate channels measure simultaneously the same DUT. Let $a(t) \leftrightarrow A(jf)$ and $b(t) \leftrightarrow B(jf)$ the background of the two instruments, and $c(t) \leftrightarrow C(jf)$ the DUT noise or any common noise. The two instrument outputs are

$$x(t) = c(t) + a(t) \quad (12.6)$$

$$y(t) = c(t) + b(t), \quad (12.7)$$

where $a(t)$, $b(t)$ and $c(t)$ are statistically independent because we have put all the common noise in $c(t)$. The cross-spectrum averaged on m measures is

$$\begin{aligned} S_{yx}(f) &= \langle YX^* \rangle_m \\ &= \langle CC^* \rangle_m + \langle CB^* \rangle_m + \langle AC^* \rangle_m + \langle AB^* \rangle_m \\ &= S_c(f) + O(\sqrt{1/m}) , \end{aligned} \quad (12.8)$$

where $O(\cdot)$ means ‘order of.’ Owing to statistical independence of $a(t)$, $b(t)$ and $c(t)$, also $A(jf)$, $B(jf)$ and $C(jf)$ are statistically independent. Hence, the cross terms decrease as $\sqrt{1/m}$. This enables to interpret $S_{yx}(f)$ as follows.

Statistical limit. With no DUT noise and with two fully independent channels, it holds that $c(t) = 0$. After (12.8), the statistical limit of the measurement is

$$S_{yx}(f) \approx \sqrt{\frac{1}{m}} S_a(f) S_b(f) \quad (\text{statistical limit}). \quad (12.9)$$

Accordingly, a 5 dB improvement on the single-channel noise costs a factor of 10 in averaging, thus in measurement time.

Correlated hardware background. Still at zero DUT noise, we break the hypothesis of statistical independence of the two channels. We interpret $c(t)$ as the correlated noise of the instrument, due to environment, to the crosstalk between the two channel, etc. This is the hardware limit of the instrument sensitivity

$$S_{yx}(f) \simeq [S_c(f)]_{\text{xtalk, etc.}} \quad (\text{hardware limit}). \quad (12.10)$$

Regular DUT measurement. Now we introduce the DUT noise. Under the assumptions that

1. m is large enough for the statistical limit to be negligible, and
 2. the single-channel background is negligible as compared to the DUT noise,
- the cross spectrum gives the *DUT noise*

$$S_{yx}(f) \simeq [S_c(f)]_{\text{DUT}} \quad (\text{DUT measurement}). \quad (12.11)$$

This is the regular use of the instrument.

12.1.3 Delay line theory

Delaying the signal $v(t)$ by τ , all time-varying parameters of $v(t)$ are also delayed by τ , thus the phase fluctuation $\varphi(t)$ turns into $\varphi(t-\tau)$. By virtue of the time-shift theorem, the Fourier transform of $\varphi(t-\tau)$ is $e^{-j2\pi\tau f}\Phi(jf)$. This enables the measurement of the oscillator phase noise $\varphi(t)$ by observing the difference $\theta(t) = \varphi(t) - \varphi(t-\tau)$. Referring to Fig. 12.1, the output signal is $V_o(jf) = k_\varphi \Theta(jf) = k_\varphi H(jf)\Phi(jf)$, where k_φ is the mixer phase-to-voltage gain, and $H(jf) = 1 - e^{-j2\pi\tau f}$ is the system transfer function. Consequently

$$S_v(f) = k_\varphi^2 |H(jf)|^2 S_\varphi(f) \quad (12.12)$$

$$|H(jf)|^2 = 4 \sin^2(\pi f \tau) . \quad (12.13)$$




Figure 12.1: Basic delay-line phase noise measurement.




Figure 12.2: Oscillator phase-noise model.

The oscillator noise $S_\varphi(f)$ is inferred by inverting (12.12). In practice, it is important to keep $S_v(f)$ accessible because it reveals most of the experimental mistakes connected with the instrument background. The function $|H(jf)|^2$ has a series of zeros at $f = n/\tau$, integer n , in the vicinity of which the experimental results are not useful. In practice, the first zero sets the maximum measurement bandwidth to $0.95/\tau$, as discussed in [?]. Nonetheless, the regions between contiguous zeros are useful diagnostics for the oscillator under test, provided the frequency resolution of the FFT analyzer be sufficient.

At low frequency, the instrument background is naturally optimized for the measurement of $1/f^2$ and $1/f^3$ noise because $1/|H(jf)|^2 \sim 1/f^2$ for $f \rightarrow 0$, and an additional factor $1/f$ comes from the electronics. This is a fortunate outcome because the noise of the oscillators of major interest is proportional to $1/f^3$ or to $1/f^4$ in this region. Of course, an appropriate choice of τ is necessary.

12.1.4 Oscillator phase noise

The oscillator consists of an amplifier of gain A (assumed constant vs. frequency) and of a feedback transfer function $\beta(jf)$ in closed loop. The gain A compensates for the losses, while $\beta(jf)$ selects the oscillation frequency. This model is general, independent of the nature of A and $\beta(jf)$. We assume that the Barkhausen condition $|A\beta(jf)| = 1$ for stationary oscillation is verified at the carrier frequency ν_0 by saturation in the amplifier or by some other gain-control mechanism. Under this hypothesis, the phase noise is modeled by the scheme of Fig. 12.2, where all signals are *the phases of the oscillator loop* [?,?]. This model is inherently linear, so it eliminates the mathematical difficulty due to the parametric nature of flicker noise and of the noise originated from the environment fluctuations. We denote with $\varphi(t) \leftrightarrow \Phi(jf)$ the

oscillator phase noise, and with $\psi(t) \leftrightarrow \Psi(jf)$ the amplifier phase noise. The latter is always additive, regardless of the underlying physical mechanism. More generally, $\psi(t)$ accounts for the phase noise of all the electronic and optical components in the loop. The ideal amplifier ‘repeats’ the phase of the input, for it has a gain of one (exact) in the phase-noise model. The feedback path is described by the transfer function $B(jf)$ of the phase perturbation. In the case of the *delay-line oscillator*, the feedback path is a delay line of delay τ_d followed by a resonator of relaxation time $\tau_f \ll \tau_d$ that selects the oscillation frequency ν_0 among the multiples of $1/\tau_d$. Neglecting the difference between natural frequency and oscillation frequency, τ_f is related to the quality factor Q by $\tau_f = \frac{Q}{\pi\nu_0}$. The phase-perturbation response of the feedback path is

$$B(jf) = \frac{e^{-j2\pi f\tau_d}}{1 + j2\pi f\tau_f} . \quad (12.14)$$

The oscillator is described by the phase noise transfer function

$$H(jf) = \frac{\Phi(jf)}{\Psi(jf)} \quad [\text{definition of } H(jf)] . \quad (12.15)$$

Using the basic equations of feedback, by inspection on Fig. 12.2 we find

$$H(jf) = \frac{1}{1 - B(jf)} , \quad (12.16)$$

and consequently

$$|H(jf)|^2 = \frac{1 + 4\pi^2 f^2 \tau_f^2}{2 - 2 \cos(2\pi f\tau_d) + 4\pi^2 f^2 \tau_f^2 + 2\omega\tau_f \sin(2\pi f\tau_d)} . \quad (12.17)$$

The detailed proof of (12.14) and (12.17) is given in [?, ?]. The result is confirmed using the phase diffusion and the formalism of stochastic processes [?]. Figure 12.3 shows an example of OEO phase-noise. The loop noise is $S_\psi(f) = 8 \times 10^{-12}/f + 10^{-14}$, the same used in Sec. 12.4.6. The periodicity of the delay-line phase produces a series of noise peaks at frequency multiple of $1/\tau_d$. These peaks have extremely narrow bandwidth, in the Hertz range, but in simulations and experiments they seem significantly wider because of the insufficient resolution. The peak height follows the law $S_\varphi(f) \sim 1/f^4$ (−40 dB/decade). Eliyahu [?, Sec.IV] reports about a discrepancy by a factor \sqrt{f} (−35 dB/decade), yet giving little details and suggesting that further investigation is necessary.

12.2 Optical-fiber microwave delay

12.2.1 Design strategy

The optical-fiber delay unit, shown in Fig. 12.4, is designed for low noise and for high stability of the delay, as discussed underneath. Focusing on the X band, for practical




Figure 12.3: Calculated OEO phase-noise.

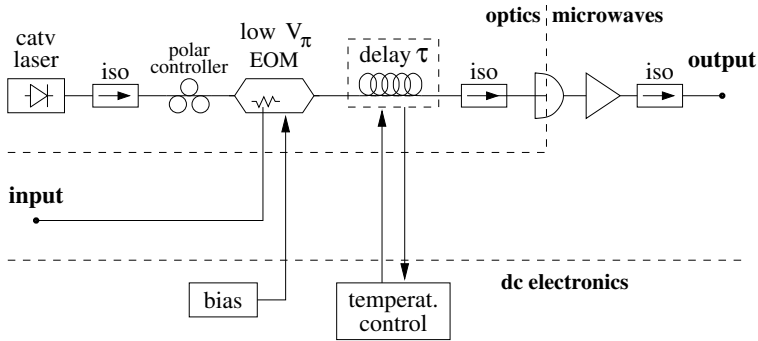


Figure 12.4: Optical-fiber delay unit.

reasons we chose standard 7–12.4 GHz microwave components, mainly intended for radar and telecommunication technology.

The optical fiber is a Corning SMF-28 wound on a cylinder of 15 cm diameter and 2 cm height. We used 2 km (10 μ s) and 4 km (20 μ s) fibers in most experiments, and sometimes shorter ones. The spool is enclosed in a 5 mm thick Duralumin cylinder thermally insulated from the environment by a 3 cm plastic foam layer. The cylinder is stabilized at room temperature within a fraction of a milliKelvin with a PID control built in our laboratory, and set with the Ziegler-Nichols method. The ultimate period, i.e., the inverse oscillation frequency of the proportional-only control at the oscillation threshold, is of 40 s. The advantage of the control vs. a passive stabilization (large metal mass and insulator) is still questionable. In the short-term the passive stabilization would certainly be preferable because it does not suffer from the noise inherent in the control. On the other hand, in phase noise measurements (Fig. 12.5) we need to stabilize the quadrature condition at the mixer inputs during the session, which last up to one day. Moreover, in dual-channel measurements (Section 12.3) the residual environment fluctuations are fully correlated, while the thermal fluctuations

of the control can be rejected because the two controls are independent.

The light source is a semiconductor CATV laser, temperature controlled and powered with a low-noise current source. This choice is partially motivated by the need of reasonable simplicity and of low RIN (relative intensity noise). The RIN turns into AM noise of the detected microwave signal, which pollutes the phase noise measurements (Section 12.3.1). During the past two years we used both $1.31\ \mu\text{m}$ and $1.55\ \mu\text{m}$ lasers. The fiber attenuation is of $0.2\ \text{dB/km}$ at $1.55\ \mu\text{m}$, and of $0.35\ \text{dB/km}$ at $1.31\ \mu\text{m}$. The dispersion of the SMF-28 fiber goes to zero at $1.311 \pm 0.01\ \mu\text{m}$, which virtually eliminates the effect of the laser frequency noise in the vicinity of that wavelength. For reference, at $1.55\ \mu\text{m}$ the dispersion is of $17\ \text{ps/nm/km}$. A laser linewidth of $10\ \text{MHz}$ ($5.8 \times 10^{-5}\ \text{nm}$) at $1.55\ \mu\text{m}$ produces a delay fluctuation of $2 \times 10^{-15}\ \text{s}$ rms after $2\ \text{km}$ optical fiber, which is equivalent to $1.2 \times 10^{-4}\ \text{rad}$ at the microwave frequency of $10\ \text{GHz}$. The frequency flicker of one of our lasers, measured in the frequency domain with an asymmetric Mach-Zehnder and converted into Allan deviation $\sigma_y(\tau)$, is of 4×10^{-10} . This preliminary result indicates that the laser fluctuation should be of less than $100\ \text{kHz}$ (flicker floor), which gives a phase-noise contribution lower than other noises. In the end, we have a weak preference for the $1.55\ \mu\text{m}$ lasers, based on the more progressed technology, and after comparing empirically the effect several lasers on the phase noise spectra.

The intensity modulator is a Mach-Zehnder electro-optic modulator (EOM) exhibiting low half-wave voltage ($V_\pi \simeq 3.9\ \text{V}$), so that the maximum modulation is achieved with no more than $50\ \text{mW}$ ($+17\ \text{dBm}$) of microwave power. This choice is important for the stability of the bias point because the LiNbO_3 is highly sensitive to temperature, thus to power and to thermal gradients. Other modulation methods have been discarded, the direct modulation of the laser because the laser threshold enhances the microwave phase noise, and the acusto-optic modulator because it is unsuitable to microwaves.

For low noise, the photodetector can only be a InGaAs p-i-n diode operated in strong reverse-bias conditions, thus a photoconductor. Reverse bias is necessary for high speed, as it reduces the capacitance. The need for low noise excludes some other detectors, like the avalanche diode. The traditional photodetectors loaded to a resistor are preferred to the more modern ones with internal transconductance amplifier because of the possibility to choose a low-flicker external amplifier.

12.2.2 Output power and white noise

Using the subscript λ for light and the overline for the time-average, the modulated optical power at the output of the EOM is

$$P_\lambda(t) = \overline{P}_\lambda [1 + m_i \cos(2\pi\nu_0 t)] , \quad (12.18)$$

where the intensity-modulation index is

$$m_i = 2J_1 \left(\frac{\pi V_p}{V_\pi} \right) , \quad (12.19)$$

$J_1()$ is the 1st order Bessel function, V_p is the microwave peak voltage, and V_π is the modulator half-wave voltage. Equation (12.19) originates from the sinusoidal nature

of the EOM response, with $V_p \cos(2\pi\nu_0 t)$ input voltage. The harmonics at frequency $n\nu_0$, integer $n \geq 2$, fall beyond the microwave bandwidth, thus they are discarded. Though the maximum modulation index is $m_i = 1.164$, occurring at $V_p = 0.586 V_\pi$, the practical values are 0.8–1.

The detected photocurrent is $i(t) = \rho P_\lambda(t)$, where ρ is the detector responsivity. Assuming a quantum efficiency of 0.6, the responsivity is of 0.75 A/W at 1.55 μm wavelength, and of 0.64 A/W at 1.31 μm . The dc component of $i(t)$ is $i_{\text{dc}} = \rho \bar{P}_\lambda$. The microwave power at the detector output is

$$P_0 = \frac{1}{2} m_i^2 R_0 \rho^2 \bar{P}_\lambda^2 \quad (\text{detector output}) , \quad (12.20)$$

where $R_0 = 50 \Omega$ is the load resistance.

The white noise at the input of the amplifier is

$$N = F k T_0 + 2q R_0 \rho \bar{P}_\lambda \quad (\text{white noise}) , \quad (12.21)$$

where F is the amplifier noise figure and $k T_0 = 4 \times 10^{-21}$ J is the thermal energy at room temperature. The first term of (12.21) is the noise of the amplifier, and the second term is the shot noise. Using $b_0 = N/P_0$, the white phase noise is

$$b_0 = \frac{2}{m_i^2} \left[\frac{F k T_0}{R_0} \frac{1}{\rho^2 \bar{P}_\lambda^2} + \frac{2q}{\rho \bar{P}_\lambda} \right] \quad (\text{white phase noise}) . \quad (12.22)$$

Interestingly, the noise floor is proportional to $1/\bar{P}_\lambda^2$ at low power, and to $1/\bar{P}_\lambda$ above the threshold power

$$P_\lambda = \frac{F k T_0}{R_0} \frac{1}{2\rho q} \quad (\text{threshold power}) . \quad (12.23)$$

For example, taking $\rho = 0.75$ A/W and $F = 5$ (SiGe parallel amplifier), we get a threshold of 1.7 mW, at which the noise floor is $b_0 \simeq 10^{-15}$ rad^2/Hz (–150 dBrad²/Hz).

12.2.3 Flicker noise

Phase and amplitude flickering result from the near-dc $1/f$ noise up-converted by non-linearity or by a parametric modulation process. This is made evident by two simple facts:

1. The $1/f$ noise is always always present in the dc bias of electronic devices.
2. In the absence of a carrier, the microwave spectrum at the output of a device is white, i.e. nearly constant in a wide frequency range.

Assuming that the phase modulation is approximately linear unless the carrier is strong enough to affect the dc bias, two basic rules hold:

1. b_{-1} is independent of the carrier power.
2. Cascading two or more devices the b_{-1} add up, regardless of the device order in the chain.

The reader may have in mind the Friis formula for the noise referred to the input of a chain [?], stating that the contribution of each stage is divided by the total gain of the preceding stages, and therefore indicating the first stage as the major noise source. The Friis formula raises from the additive property of white noise, hence it does not apply to parametric noise.

Early measurements on *amplifiers* [?,?,?] suggest that different amplifiers based on a given technology tend to have about the same b_{-1} , and that b_{-1} is nearly constant in a wide range of carrier frequency and power. Our independent experiments confirm that the $1/f$ phase noise of a given amplifier is independent of power in a wide range [?], [?], [?, Chapter 2]. For example, b_{-1} of a commercial amplifier (Microwave Solutions MSH6545502) that we measured at 9.9 GHz is between 1.25×10^{-11} and 2×10^{-11} from 300 μW to 80 mW of output power. Similarly, the $1/f$ noise of a LNPT32 SiGe amplifier measured between 32 μW (−15 dBm) and 1 mW (0 dBm) in 5 dB steps overlap perfectly. In summary, the typical phase flickering b_{-1} of a “good” microwave amplifier is between 10^{-10} and 10^{-12} rad^2/Hz (−100 to −120 dBrad²/Hz) for the GaAs HBTs, and between 10^{-12} and 10^{-13} rad^2/Hz (−120 to −130 dBrad²/Hz) for the SiGe transistors.

The $1/f$ noise of the *microwave photodetector* is expected to be similar to that of an amplifier because the underlying physics and technology are similar. The measurement is a challenging experimental problem, which has been tackled only at the JPL independently by Rubiola and Shieh [?,?,?]. The results agree in that at 10 GHz the typical b_{-1} of a InGaAs p-i-n photodetector is of 10^{-12} rad^2/Hz (−120 dBrad²/Hz).

Microwave *variable attenuators* and *variable phase shifters* can be necessary for adjustment. Our early measurements [?] indicate that the b_{-1} of these components is of the order of 10^{-15} rad^2/Hz (−150 dBrad²/Hz), which is negligible as compared to the amplifiers and to the photodetectors.

Additional sources of noise are the EOM, the laser amplified spontaneous emission, and the noise of the optical pump. As theory provide no indications about these effects, a pragmatic approach is necessary, which consists of measuring the total noise of the microwave delay unit in different configurations.

12.3 Dual-delay phase noise measurement

Figure 12.5 shows the scheme of the instrument. However similar to the previous ones ([?,?]), this version adds engineering and a substantial progress in understanding $1/f$ noise.

The instrument consists of two equal and independent channels that measure the oscillator under test using the principle of Fig. 12.1. Then, the single-channel noise is removed using the cross-spectrum method before converting the mixer output into the oscillator noise $S_\varphi(f)$ with Eq. (12.13).

Looking at one channel, we observe that the microwave signal is split into two branches before the EOM, so that the long branch consists of modulator, optical fiber (delay τ), photodetector and microwave amplifier, while the short branch is a microwave path of negligible length. This differs from the first single-channel instrument [?, Fig. 7], in which the short branch was optical. Removing the photodetector

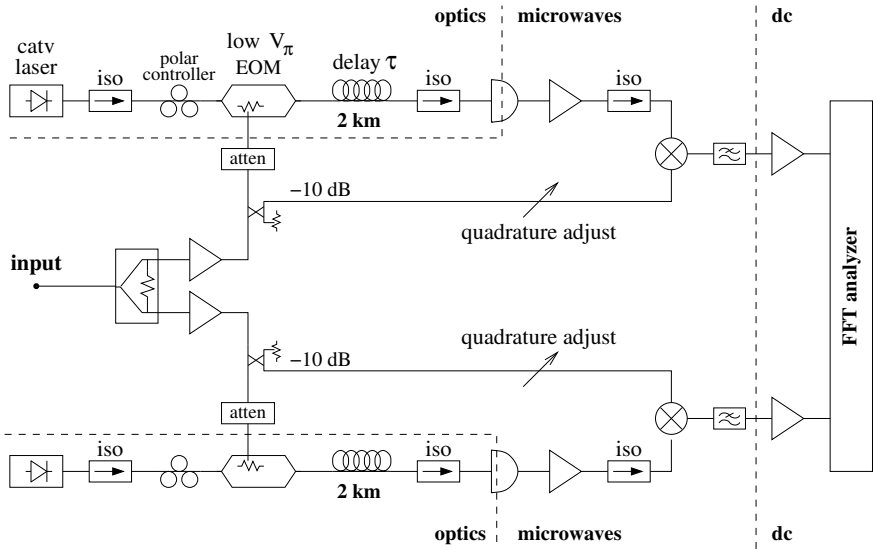


Figure 12.5: Scheme of the dual-channel instrument.

and the microwave amplifier from the short branch yields lower noise, and in turn faster convergence of the correlation algorithm. Lower laser power is needed. Another advantage is that we use a microwave power splitter instead of an optical power splitter. While the noise of the former is negligible for our purposes [?, ?], we have no first-hand knowledge about the latter. The problem with this configuration is that we no longer have an optical input, so the microwave-modulated light beams can not be measured.

Trading off with the available components, we had to use both GaAs amplifiers and SiGe amplifiers. Since the SiGe amplifiers exhibit lower $1/f$ noise, they are inserted at the photodetector output. This choice is motivated by the fact that the $1/f$ noise at the mixer input is converted into oscillator $1/f^3$ noise by Eq. (12.12)–(12.13) while the $1/f$ noise at the EOM input is not.

12.3.1 Mixer Noise

Used as a phase detector, the double-balanced mixer needs to be saturated at both inputs. The conversion gain is of 0.1–0.5 V/rad when the power is in the appropriate range, which is of ± 5 dB centered around the optimum power of 5–10 mW. Out of this range, b_{-1} increases. At lower power the conversion gain drops suddenly because the input voltage is insufficient for the internal Schottky diodes to switch.

Out of experience in low-flicker applications, we have a preference for the mixers manufactured by Narda and by Marki. The coefficient b_{-1} is of the order of 10^{-12} , similar to that of the photodetectors. The white noise is chiefly the noise of the output amplifier divided by the conversion gain k_φ . Assuming that the amplifier noise is 1.6 nV/ $\sqrt{\text{Hz}}$ (our low-flicker amplifiers input-terminated to 50 Ω [?]) and that $k_\varphi = 0.1$

V/rad (conservative with respect to P_0), the white noise is $b_0 = 2.5 \times 10^{-16} \text{ rad}^2/\text{Hz}$ ($-156 \text{ dBrad}^2/\text{Hz}$).

Mixers are sensitive to the amplitude noise $\alpha(t)$ of the input signal, hence the mixer output takes the form $v(t) = k_\varphi \varphi(t) + k_\alpha \alpha(t)$. This is due to the asymmetry of the internal diodes and baluns. In some cases we have measured k_φ/k_α of 5 or less, while values of 10–20 are common. In photonic systems the contamination from amplitude noise can be a serious problem because of the power of some lasers and laser amplifiers fluctuates. Brendel [?] and Cibel [?] suggest that the mixer can be operated at a sweet point off the quadrature, where the sensitivity AM noise nulls. A further study shows that the Brendel offset method can not be used in our case [?] because the null of amplitude sensitivity results from the equilibrium between equal and opposite sensitivities at the two inputs. But the delay de-correlates the mixer input signals by virtue of the same mechanism used to measure $S_\varphi(f)$.

12.3.2 Calibration

The phase-noise measurement system is governed by $S_v(f) = k_\varphi^2 |H(jf)|^2 S_\varphi(f)$ [Eq. (12.12) and (12.13)]. Calibration takes the accurate measurement of τ and k_φ . An accuracy of 1 dB can be expected. Since the stability of the optical fiber exceeds our needs, τ does not need re-calibration after first set-up. Optical reflection methods are suitable, as well as the spectrum measurement of the noise the peak at $f = 1/\tau$, where $|H(jf)|^2 = 0$. Conversely, k_φ is highly sensitive to optical and microwave power. It is therefore recommended to measure it at least every time the experimental conditions are changed.

The simplest way to measure k_φ is to introduce a reference sinusoidal modulation $\varphi(t) = m_\varphi \sin 2\pi f_m t$ in the microwave signal (12.1), after replacing the oscillator under test with a microwave synthesizer. This accounts for the dc amplifier at the mixer output, not shown in Fig. 12.1. It is convenient to set $f_m \lesssim 0.1/\tau$, so that it holds $\sin(\pi f_m \tau) \simeq \pi f_m \tau$ in (12.13). With commercial synthesizers, frequency modulation is more suitable than phase modulation because the frequency deviation Δf can be measured in static conditions, with dc input. Frequency modulation is equivalent to phase modulation

$$\varphi(t) = \frac{\Delta f}{f_m} \sin 2\pi f_m t \quad (12.24)$$

of index $m_\varphi = \Delta f/f_m$. Substituting (12.24) in (12.1) and truncating the series expansion to the first-order, the microwave signal is

$$v(t) = J_0\left(\frac{\Delta f}{f_m}\right) \cos 2\pi \nu_0 t + J_1\left(\frac{\Delta f}{f_m}\right) [\cos 2\pi(\nu_0 + f_m)t - \cos 2\pi(\nu_0 - f_m)t] , \quad (12.25)$$

where $J_n(\cdot)$ is the Bessel function of order n . For small m_φ , we use the approximations $J_0(m_\varphi) \simeq 1$ and $J_1(m_\varphi) \simeq \frac{1}{2} m_\varphi$. Thus,

$$v(t) = \cos 2\pi \nu_0 t + \frac{1}{2} \frac{\Delta f}{f_m} [\cos 2\pi(\nu_0 + f_m)t - \cos 2\pi(\nu_0 - f_m)t] . \quad (12.26)$$




Figure 12.6: Measured single-channel background noise with zero-length optical fiber.

Finally, the modulation index can be easily calibrated by inspection with a microwave spectrum analyzer. In this case it is recommended to measure m_φ with relatively large sidebands (say, 40 dB below the carrier), and then to reduce m_φ by inserting a calibrated attenuator at the FM input of the synthesizer. This enhances the accuracy of the spectrum analyzer. For example, having $\tau = 10 \mu\text{s}$ and $k_\varphi = 0.2 \text{ V/rad}$, we may set $f_m = 5 \text{ kHz}$ and $m_\varphi = 2 \times 10^{-2} \text{ rad}$. In this case, the detected signal has a peak voltage of $400 \mu\text{V}$ at the mixer output, thus 40 mV after 40 dB amplification. Measuring k_φ , it may be convenient to attenuate the modulating signal by 20–30 dB, so that the system is calibrated in actual operating conditions.

12.4 Experiments

12.4.1 Measurement of the background at zero fiber length

Replacing the spools of Fig. 12.5 with short fibers, the oscillator phase noise is rejected. The noise phenomena originated inside the fiber, or taken in by the fiber, are also eliminated. The noise measured in this conditions (Fig. 12.6) is the instrument background as it would be with noise-free fibers. The curve a (red) is the phase noise $S_\theta(f)$ measured the mixer output. The other curves (b–f) are plotted using the same data set, after unapplying $k_\varphi^2 |H(jf)|^2$ for various fiber lengths [Eq. (12.12)–(12.13)]. As expected after Sec. 12.1.3, the curve f (black, 4 km fiber, $\tau = 20 \mu\text{s}$) shows a peak at $1/\tau = 50 \text{ kHz}$ and at $2/\tau = 100 \text{ kHz}$, where $|H(jf)|^2 = 0$. Additionally, the curve f shows two minima 6 dB lower than $S_\theta(f)$ at $f = 25 \text{ kHz}$ and at $f = 75 \text{ kHz}$, where $|H(jf)|^2 = 4$. The same facts are observed at twice the frequency on the curve e (blue, 2 km fiber, $\tau = 10 \mu\text{s}$).




Figure 12.7: Effect of the laser RIN and of the oscillator-under-test AM noise, measured with zero-length fiber.

12.4.2 Effect of microwave AM noise and of laser RIN

This experiment shows qualitatively the effect of the laser RIN and of the AM noise of the oscillator under test, still in single-channel mode and with zero-length optical fiber, so that the oscillator phase noise is rejected. The curve 3 (red) is the same as the curve a of Fig. 12.6. In this case, the optical source is a CATV laser, while the oscillator under test is a 10 GHz sapphire-loaded dielectric-cavity oscillator operated at room temperature [?]. Besides high stability, the sapphire oscillator performs low AM noise. Replacing the laser with a different one with higher RIN (curve 2, blue), or replacing the oscillator under test with a synthesizer (curve 3, black), which has higher AM noise, the background increases significantly.

12.4.3 Measurement of a sapphire microwave oscillator

We measured the phase-noise of a room-temperature sapphire oscillator, still in single-channel mode, progressively increasing the fiber length. Let us focus on the $1/f^3$ noise, which dominates on the spectrum (Fig. 12.8). When the delay is insufficient to detect the noise of the source under test, the spectrum is the background of the instrument, which scales with the inverse square length, that is, -6 dB in $S_\varphi(f)$ for a factor of 2 in the delay. This is visible on the curves a (magenta, 100 m, $\tau = 0.5 \mu\text{s}$) and b (red, 500 m, $\tau = 2.5 \mu\text{s}$). Increasing the length, the $1/f^3$ noise no longer decreases. This fact, seen on the curves d (blue, 2 km, $\tau = 10 \mu\text{s}$) and e (black, 4 km, $\tau = 20 \mu\text{s}$), indicates that the instrument measures the phase noise of the sapphire oscillator, which does not scale with the delay. Measuring the $1/f^3$ noise of a sapphire oscillator in single-channel mode is a remarkable result because this oscillator is regarded as the best reference in the field of low-noise microwave sources.




Figure 12.8: Phase noise of a sapphire oscillator.

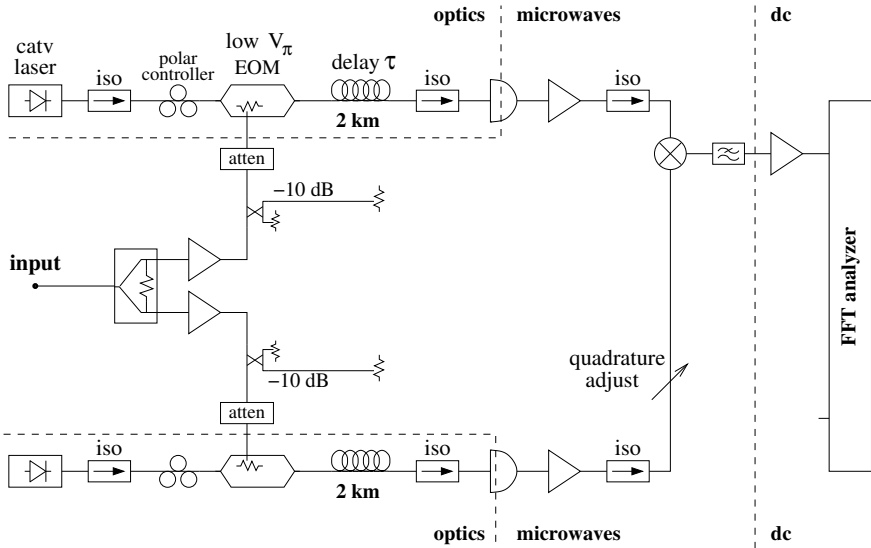


Figure 12.9: Measurement of the background noise, including the optical fibers.

12.4.4 Assessing the noise of the photonic channel

We measured the phase noise of the photonic channel using the scheme of Fig. 12.9, derived from Fig. 12.5 after removing some parts. The phase noise of the reference oscillator is rejected by using two fibers of equal length. This experiments suffers from the following limitations.

1. The noise of the fibers can not be separated from other noises.




Figure 12.10: Background noise, including the optical fibers.

2. The $1/f$ noise of the GaAs amplifiers that drive the EOM can show up.
3. We could not use the correlation method because it takes four matched optical delay-lines, which were not available.

Nonetheless, this scheme has the merit of giving at least an upper bound of the achievable noise. The measured spectrum, shown in Fig. 12.10, indicates that the $1/f$ phase noise is $b_{-1} = 8 \times 10^{-12} \text{ rad}^2/\text{Hz}$ ($-111 \text{ rad}^2/\text{Hz}$). At this level, the mixer noise is negligible.

12.4.5 Background noise of the two-channel instrument

We measured the background noise in the two-channel configuration, using the cross-spectrum method of Sec. 12.1.2 and with zero-length optical fiber, so that the phase noise of the 10 GHz reference oscillator is rejected. This experiment does not account for the optical noises originated inside the fibers, these phenomena are rejected in Eq. (12.8) because the two fibers can not be correlated. When this experiment was done, the stability of the quadrature condition was still insufficient for long acquisitions. For this reason we stopped the measurement after $m = 200$ spectra. The cross spectrum is shown in Fig. 12.11.

The reference straight line (red) is the $1/f^3$ phase noise equivalent to the Allan deviation $\sigma_y = 10^{-12}$, calculated with Eq. (12.4) and (12.5). Thus, averaging on 200 spectra the instrument can measure the stability of an oscillator at 10^{-12} level.

The signature of a correlated noise is a smooth cross spectrum. Yet, in Fig. 12.11 the variance is far too large for the two channels to be correlated. This indicates that the sensitivity is limited by the small number of averaged spectra. Additionally, the spectrum shows a sawtooth-like discrepancy vs. the $1/f^3$ line. This averaging effect, due to the measurement bandwidth that increases with frequency in logarithmic




Figure 12.11: Background noise in two-channel mode, measured with zero-length fiber.




Figure 12.12: Scheme of the opto-electronic oscillator.

resolution, further indicates that the value of m is still insufficient. The ultimate sensitivity for large m is still not known.

12.4.6 Opto-electronic oscillator

We implemented the oscillator of Fig. 12.12 using a 4 km delay line (20 μ s), a SiGe amplifier, and a photodetector with integrated transconductance amplifier. Oscillation starts at 6–7 mW optical power. The best working point occurs at 12–13 mW optical power, where the microwave output power is 20 mW. For flicker noise the microwave chain is similar to one channel in Fig. 12.9 because the order of the devices in the chain is not relevant. Thus, we take $b_{-1} = 8 \times 10^{-12} \text{ rad}^2/\text{Hz}$ (–111 dB rad²/Hz), the same of Fig. 12.10, as the first estimate of the loop noise. Though somewhat arbitrary, this value accounts for the photodetector internal amplifier, more noisy than our microwave amplifiers. Using $(b_{-1})_{\text{loop}} = 8 \times 10^{-12} \text{ rad}^2/\text{Hz}$ in the oscillator noise model of Section 12.1.4, the expected oscillator flickering is $(b_{-3})_{\text{osc}} = 6.3 \times 10^{-4} \text{ rad}^2/\text{Hz}$




Figure 12.13: Phase noise of the opto-electronic oscillator.

($-32 \text{ dBrad}^2/\text{Hz}$). By virtue of (12.4) and (12.5), this is equivalent to a frequency stability $\sigma_y = 2.9 \times 10^{-12}$ (Allan deviation). The noise spectrum (Fig 12.13), measured with the dual-channel instrument, shows a frequency flicker of $10^{-3} \text{ rad}^2/\text{Hz}$ ($-30 \text{ dBrad}^2/\text{Hz}$), equivalent to $\sigma_y = 3.6 \times 10^{-12}$. The discrepancy between predicted value and result is of 2 dB.

Interestingly, the OEO phase noise compares favorably to the lowest-noise microwave synthesizers and quartz oscillator multiplied to 10 GHz. Yet, the OEO can be switched in steps of 50 kHz without degrading the noise.

12.5 Further developments

Our experiments suggest the following improvements, tested or in progress.

1. In the phase-noise instrument, the microwave power the input of the EOM and the two inputs of the mixer is a critical parameter because the gain of the entire system is strongly non-linear. The insertion of test points is recommended, using common and inexpensive power detectors after tapping a fraction of the power with 20 dB directional couplers.
2. The instrument requires that the two signals at the mixer input are kept in quadrature. This is usually accomplished with a line stretcher or a with voltage-controlled phase shifter. A smarter solution exploits the dispersion of the optical fiber, by adjusting the laser wavelength via the temperature control [?]. This is viable only at $1.55 \mu\text{m}$ and with long fibers. For reference, using a 2 km fiber and our lasers, it takes 10 K of temperature change for a quarter wavelength of the microwave signal. Another solution exploits the temperature coefficient of the fiber delay, $6.85 \times 10^{-6}/\text{K}$. Accordingly, a quarter wavelength takes 0.37 K of temperature with 2 km fiber (10 μs), or 3.7 K with 200 m.

3. Some popular EOMs have a low-frequency photodetector at the unused output port of the Mach-Zehnder interferometer, which we hope to exploit to stabilize the bias point.
4. The OEO frequency can be fine-tuned by adding a RF signal with an SSB modulator at the output. Extremely high resolution, of the order of 10^{-16} can be obtained with a 48 bit DDS, with no degradation of the spectral purity [?, ?].

Acknowledgments

We are indebted with Lute Maleki (OEwaves and NASA/Caltech JPL, Pasadena, CA, USA), Nan Yu (NASA/Caltech JPL) and Ertan Salik (NASA/Caltech JPL, now with the California State Polytechnic University, Pomona, CA, USA) for having taught us a lot in this domain, and for numerous discussions, advices and exchanges. Gilles Cibiel (CNES, France) provided contract support and help. Rodolphe Boudot (FEMTO-ST, now with SYRTE, Paris, FRANCE) helped with the sapphire oscillator, Xavier Jouvenceau (FEMTO-ST) helped in the implementation of the early version of the correlation system, and Cyrus Rocher (FEMTO-ST) implemented the temperature control.

This work is supported with grants from Aeroflex, CNES and ANR.

12.6 Photodetector

File pd.tex³

A photodiode is a diode whose junction can be illuminated through an aperture or an optical fiber, and optimized for the detection of light. Any diode (or bipolar transistor) is transformed into a photodiode by opening the case (modern and old metal-case device), or by removing the opaque painting (earliest glass-case devices), if the junction is made accessible to light.

The diode forward current i_F is related to the forward voltage v_F (Fig. 12.14) by

$$i_F = I_s \left[e^{\frac{v_F}{nV_T}} - 1 \right] - i_P \quad (12.27)$$

where I_s is the inverse saturation current, n is a technical parameter⁴ ($1 \leq n \leq 2$), $V_T = kT/q$ is the thermal voltage ($V_T \simeq 25.6$ mV at room temperature), and i_P is the photocurrent.

Let us first consider the photodiode in the dark. With $i_P = 0$, Eq. (12.27) describes a bare diode. When the diode is forward biased ($v_F \gg V_T$), the term “1” is negligible as compared to the exponential. At large current the exponential is limited by the bulk resistance. In reverse bias conditions, with $v_F \ll -V_T$, the exponential is negligible and the reverse current ($i_R = -i_F$) saturates to I_s . Large reverse bias results in

³From file pd.tex

⁴The symbol η generally adopted in the literature is replaced with n to avoid confusion with the quantum efficiency, denoted with η in optics.




Figure 12.14

breakdown, which is avalanche effect. The condition in which v_F is of the order of some V_T , either positive or negative, is seldom encountered in practice.

When a photon of energy E greater than the chemical potential E_c of the lattice is absorbed, it creates an electron-hole pair free to propagate. The hole-electron pair carries a charge equivalent to q . This is stated by the Ramo theorem. The physical meaning of this theorem is that a single carrier, either electron or hole, yields a triangular current pulse, and that two triangles are needed for the integral to be equal to q . A flux Φ of photons produces a photocurrent

$$i_P = q\eta\Phi, \quad (12.28)$$

where η is the quantum efficiency. η accounts for (1) the probability absorbing a photon and turning it into an electron-hole pair, (2) the back reflection at the surface, (3) the transparence, for photons pass through the lattice and are lost at the other end.

The photodiode can be operated in the four modes, shorted photovoltaic cell, loaded photovoltaic cell, photoconductor, and avalanche photoconductor.

Shorted photovoltaic cell

The terminals are shorted, or equivalently the diode is connected to the virtual ground of an operational amplifier (Fig. 12.15). This is the typical configuration for dc operation (light detectors) and for low speed applications. The photo-induced current i_P amplified produces an output voltage $v_O = -R_F i_P = R_F \eta \Phi$.




Figure 12.15: The photodiode is shorted to the virtual ground of an operational amplifier.




Figure 12.16

Loaded photovoltaic cell

The photodiode is used to convert light into electrical power delivered to a load. The power delivered is $P = -v_F i_F$. Referring to Fig. 12.14, the sign ‘-’ derives from the fact that the definition of i_F and v_F follows the convention of the *loads* ($i_F > 0$ enters into the device), while a positive delivered power requires a positive *outgoing* current. The typical load is a switching power supply that adapts itself to the illumination conditions. The maximum available power is the maximum product $P = -v_F i_F$. If lower power is needed, the working point should go towards the v_F axis (higher voltage and lower current). This is necessary because the switching power supply (and electronic circuits in general) work better at low current than at low voltage.

Photoconductor

The photodetector is connected as in Fig. 12.16 and strongly reverse-biased, yet far from the breakdown voltage. The load line of the generator shown in Fig. 12.14 is in the III quadrant instead of in the I quadrant, which is the usual location. This occurs because the diode is inserted upside down as the load in Fig. 12.16, with the positive voltage and current that are $-v_F$ and $-i_F$, respectively.

Avalanche photoconductor

12.7 Common mistakes




Figure 12.17

Figure 12.18: Energy levels in a *p-i-n* photodiode.

12.7.1 Overvoltage

12.7.2 Bias circuit

12.8 High-speed photodetectors

12.9 Method and choice of the measurement scheme

12.10 Background noise

12.11 Experimental problems

12.12 Results

12.13 Photodetector

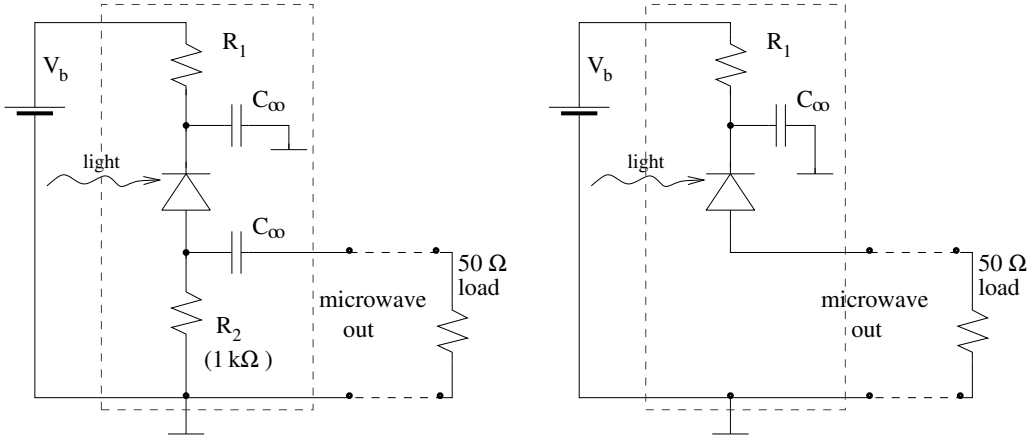


Figure 12.19

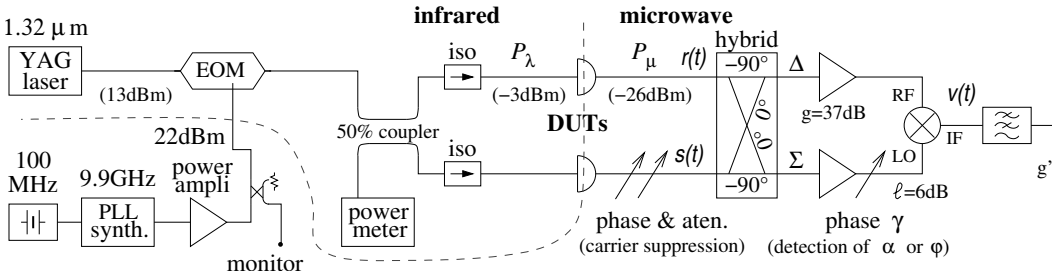


Figure 12.20

Noise in photodetectors⁵

The microwave signal at the output of a photodiode that detects a modulated optical beam contains the phase noise $\phi(t)$ and the amplitude noise $\alpha(t)$ of the detector. Beside the white noise, which is well understood, the spectral densities $S_\phi(f)$ and $S_\alpha(f)$ show flicker noise, proportional to $1/f$. We report on the measurement of the phase and amplitude noise of high-speed *p-i-n* photodiodes. The main result is that the flicker coefficient of the samples is $\sim 10^{-12} \text{ rad}^2/\text{Hz}$ ($-120 \text{ dBrad}^2/\text{Hz}$) for phase noise, and $\sim 10^{-12} \text{ Hz}^{-1}$ (-120 dB) for amplitude noise. These values could be observed only after solving a number of experimental problems and in a protected environment. By contrast, in ordinary conditions insufficient EMI isolation, and also insufficient mechanical isolation, are responsible for additional noise to be taken in. This suggests that if package and EMC are revisited, applications can take the full benefit from the surprisingly low noise of the *p-i-n* photodiodes.

⁵From IEEE T MTT 2006

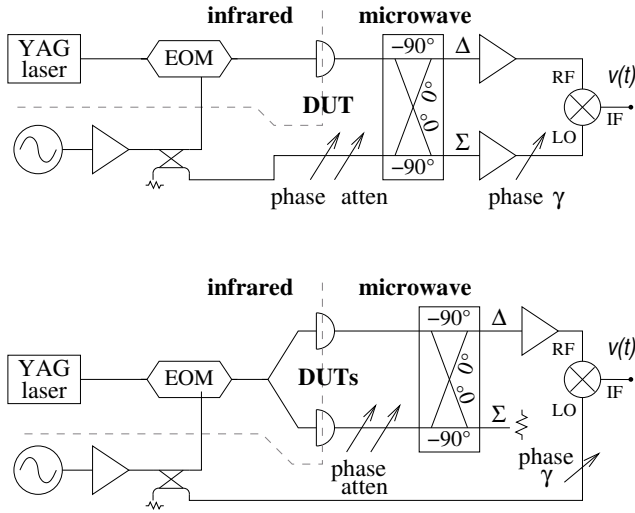


Figure 12.21: .

12.14 Introduction

Many high performance applications of microwave photonics and optics are impacted by phase noise of the microwave signals modulated as sidebands on the optical carrier. Examples of such applications include the frequency distribution system in the NASA Deep Space Network [?], very long baseline radio astronomy interferometry arrays (VLBI) [?], laboratory time and frequency comparisons [Narbonneau et al., 2003, ?], photonic oscillators [?, ?], optical microwave synthesis [?, ?, ?], and laser metrology [?]. The contributions of nearly all microwave and photonic circuit elements to the phase noise is, for most part, well understood, or at least determined experimentally. This is not the case for the contributions of the photodetector to the phase noise. Many high performance systems such as those mentioned above could be limited by the close-in noise of the photodetector. The lack of information regarding this topic made this work necessary. In fact only one conference article [?] reports on the photodiode $1/f$ noise at microwave frequencies (X band). Another work [?] was carried on at higher microwave power (10 dBm) and significantly lower frequency (1 GHz), where experiments are made simple by the availability of a commercial instrument for the direct measurement of $S_{\alpha}(f)$ and $S_{\phi}(f)$ with sufficient sensitivity. In this article we describe a sensitive measurement technique for the close-in phase noise and amplitude noise, and the measurement of several photodetectors used to detect microwave sidebands (10 GHz) of optical carriers.

When a light beam is modulated in intensity by a microwave signal and fed into a photodetector, the detector delivers a copy of the microwave signal at its output, with added noise. Flicker noise is the random fluctuation of the microwave phase and of the fractional amplitude, $\phi(t)$ and $\alpha(t)$, with power spectrum density $S(f)$ proportional

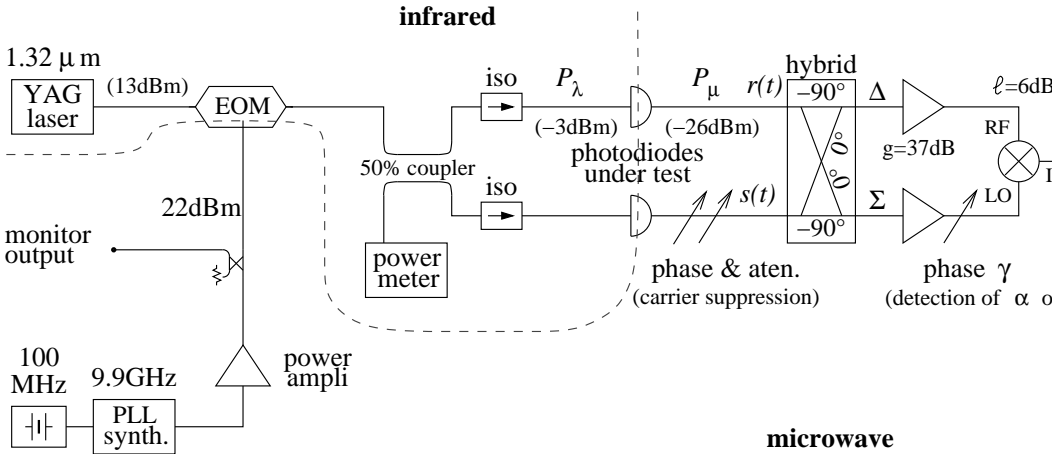


Figure 12.22: Scheme of the measurement system.

to $1/f$. This refers to the representation

$$s(t) = V_0[1 + \alpha(t)] \cos[2\pi\nu_\mu t + \phi(t)] . \quad (12.29)$$

The phase noise spectrum $S_\phi(f)$ is of paramount importance because ϕ is related to time, which is the most precisely measured physical quantity. For a review on phase noise see the References [?, ?, ?].

Most high-speed photodetectors are InGaAs *p-i-n* diodes operated in strong reverse-bias condition, hence as photoconductors. Reverse biasing is necessary for high speed because the high electric field reduces the transit time of the carriers, and also limits the junction capacitance. Thus, the depletion region (the intrinsic layer) can be tailored for quantum efficiency and speed. The *p-i-n* diode has the interesting property that even at relatively low reverse bias V_b (~ 5 V) the junction capacitance is chiefly determined by the thickness of the *i* layer [?, pp.118–119], with little influence from V_b . This indicates that phase noise may be lower than in other microwave devices.

12.15 Experimental method

A preliminary survey of the available detectors shows that none provides output power sufficient to use a saturated mixer as the phase detector, and that typical photodetectors have lower noise than common microwave amplifiers. Hence we opt for the bridge (interferometric) method, which permits flicker-free amplification before detection. This method, inspired to [?], is now a well established technique. The full theory and an extensive description of the experimental aspects is available in [?]. Hence, the description given here focus on the adaptation of the bridge method to the measurement of the photodiodes.

In our configuration (Fig.12.22) the two detector outputs are combined with ap-

propriate phase and amplitude, so that the sum (Σ) and the difference (Δ) are available at the output of the hybrid junction. At the equilibrium condition all of the microwave power goes in Σ , while only the imbalance signal, i.e., the photodetector noise plus some residual carrier, is present in Δ . Close-in flicker noise in amplifiers is a parametric effect that results from the flicker fluctuation of the dc bias that modulates the microwave carrier. Of course, the microwave output spectrum is white at zero or very low power. Hence the noise sidebands present in Δ are amplified without adding flicker. The Σ amplifier provides the power needed to saturate the LO port of the mixer, for it flickers. Yet it is shown in [?] that the close-in flickering of this amplifier is not detected because there is no carrier power on the other side of the mixer.

The detected signal, converted to dc by the mixer, is

$$v(t) = k_d \cos(\gamma + \psi) \alpha(t) - k_d \sin(\gamma + \psi) \phi(t) , \quad (12.30)$$

where ψ is the arbitrary phase that results from the circuit layout. Thus, the detection of amplitude or phase noise is selected by setting the value of γ . A fast Fourier transform (FFT) analyzer measures the output spectrum, $S_\phi(f)$ or $S_\alpha(f)$. The gain, defined as $k_d = v/\alpha$ or $k_d = v/\phi$, is

$$k_d = \sqrt{\frac{gP_\mu R_0}{\ell}} - \left[\begin{array}{c} \text{dissipative} \\ \text{loss} \end{array} \right] , \quad (12.31)$$

where g is the amplifier gain, P_μ the microwave power, $R_0 = 50 \Omega$ the characteristic resistance, and ℓ the mixer ssb loss. Under the conditions of our setup (see below) the gain is 43 dBV/[rad], including the dc preamplifier. The notation [rad] means that /rad appears when appropriate.

Calibration involves the assessment of k_d and the adjustment of γ . The gain is measured through the carrier power at the diode output, obtained as the power at the mixer RF port when only one detector is present (no carrier suppression takes place) divided by the detector-to-mixer gain. This measurement relies on a power meter and on a network analyzer. The detection angle γ is first set by inserting a reference phase modulator in series with one detector, and nulling the output by inspection with a lock-in amplifier. Under this condition the system detects α . After adding a reference 90° to γ , based either on a network analyzer or on the calibration of the phase shifter, the system detects ϕ . The phase modulator is subsequently removed to achieve a higher sensitivity in the final measurements. Removing the modulator is possible and free from errors because the phase relationship at the mixer inputs is rigidly determined by the carrier suppression in Δ , which exhibits the accuracy of a null measurement.

The background white noise results from thermal and shot noise. The thermal noise contribution is

$$S_{\phi t} = S_{\alpha t} = \frac{2FkT_0}{P_\mu} + \left[\begin{array}{c} \text{dissipative} \\ \text{loss} \end{array} \right] , \quad (12.32)$$

where F is the noise figure of the Δ amplifier, and $kT_0 \simeq 4 \times 10^{-21}$ J is the thermal energy at room temperature. This is proved by dividing the voltage spectrum $S_v =$

$\frac{2}{\ell}gFkT_0$ detected when the Δ amplifier is input-terminated, by the square gain k_d^2 . The shot noise contribution of each detector is

$$S_{\phi s} = S_{\alpha s} = \frac{4q}{\rho m^2 \overline{P}_\lambda}, \quad (12.33)$$

where q is the electron charge, ρ is the detector responsivity, m the index of intensity modulation, and \overline{P}_λ the average optical power. This is proved by dividing the spectrum density $S_i = 2q\bar{i} = 2q\rho\overline{P}_\lambda$ of the the output current i by the average square microwave current $\bar{i}_{ac}^2 = \rho^2 \overline{P}_\lambda^2 \frac{1}{2} m^2$. The background amplitude and phase white noise take the same value because they result from additive random processes, and because the instrument gain k_d is the same. The residual flicker noise is to be determined experimentally.

The differential delay of the two branches of the bridge is kept small enough (nanoseconds) so that a discriminator effect does not take place. With this conditions, the phase noise of the microwave source and of the electro-optic modulator (EOM) is rejected. The amplitude noise of the source is rejected to the same degree of the carrier attenuation in Δ , as results from the general properties of the balanced bridge. This rejection applies to amplitude noise and to the laser relative intensity noise (RIN).

The power of the microwave source is set for the maximum modulation index m , which is the Bessel function $J_1(\cdot)$ that results from the sinusoidal response of the EOM. This choice also provides increased rejection of the amplitude noise of the microwave source. The sinusoidal response of the EOM results in harmonic distortion, mainly of odd order; however, these harmonics are out of the system bandwidth. The photodetectors are operated with some 0.5 mW input power, which is low enough for the detectors to operate in a linear regime. This makes possible a high carrier suppression (50–60 dB) in Δ , which is stable for the duration of the measurement (half an hour), and also provides a high rejection of the laser RIN and of the noise of the Δ amplifier. The coherence length of the YAG laser used in our experiment is about 1 km, and all optical signals in the system are highly coherent.

12.16 Results

The background noise of the instrument is measured in two steps. A first value is measured by replacing the photodetectors output with two microwave signals of the same power, derived from the main source. The noise of the source is rejected by the bridge measurement. A more subtle mechanism, which is not detected by the first measurement, is due to the fluctuation of the mixer offset voltage induced by the fluctuation of the LO power [?]. This effect is measured in a second test, by restoring the photodetectors and breaking the path from the hybrid junction to the Δ amplifier, and terminating the two free ends. The worst case is used as the background noise. The background thereby obtained places an upper bound for the $1/f$ noise, yet hides the shot noise. This is correct because the shot noise arises in the photodiodes, not in the instrument. The design criteria of Sec. 12.15 result in a background flicker of approximately $-135 \text{ dB[rad}^2\text{]/Hz}$ at $f = 1 \text{ Hz}$, hardly visible above 10 Hz (Fig. 12.23). The white noise, about $-140 \text{ dB[rad}^2\text{]/Hz}$, is close to the expected value,




Figure 12.23: Example of measured spectra $S_\alpha(f)$ and $S_\phi(f)$.

within a fraction of a decibel. It is used only as a diagnostic check, to validate the calibration.

We tested three photodetectors, a Fermionics HSD30, a Discovery Semiconductors DSC30-1k, and a Lasertron QDMH3. These devices are InGaAs *p-i-n* photodiodes suitable to the wavelength of $1.3 \mu\text{m}$ and $1.55 \mu\text{m}$, exhibiting a bandwidth in excess of 12 GHz, and similar to one another. They are routinely used in our photonic oscillators [?, ?] and in related experiments.

Each measurement was repeated numerous times with different averaging samples in order to detect any degradation from low-frequency or non-stationary phenomena, if present. The experiment ends up with a small set of trusted spectra, each of which is the sum of the noise of two different photodiodes. Figure 12.23 shows an example. As expected, $S_\alpha(f)$ and $S_\phi(f)$ are overlapped in the white region, where noise is additive [Equations (12.32) and (12.33)]. The flicker of each device (Table 12.1) results from the solution of a linear system based on the experimental data. Each experimental spectrum is affected by a random uncertainty of 0.5 dB, due to parametric spectral estimation (Ref. [?], chap. 9), and to the measurement of the photodetector output power. In addition, we account for a systematic uncertainty of 1 dB due to the calibration of the instrument gain. The systematic uncertainty is a constant error that scales up or down all spectra by the same amount, for it also applies as is to the noise of each photodiode. Conversely, the random uncertainty increases in the process calculating the noise of the single detector, depending on the coefficients of the linear system. The practical consequence is that the fractional uncertainty becomes quite high if the device noise is lower than the average, as it happens with the HSD-30.

Table 12.1: Flicker noise of the photodiodes.

photodiode	$S_\alpha(1\text{Hz})$		$S_\phi(1\text{Hz})$	
	estimate	uncertainty	estimate	uncertainty
HSD30	-122.7	-7.1 +3.4	-127.6	-8.6 +3.6
DSC30-1K	-119.8	-3.1 +2.4	-120.8	-1.8 +1.7
QDMH3	-114.3	-1.5 +1.4	-120.2	-1.7 +1.6
unit	dB/Hz	dB	dBrad ² /Hz	dB

12.17 Discussion

For practical reasons, we selected the configurations that give reproducible spectra with low and smooth $1/f$ noise that are not influenced by the sample averaging size. Reproducibility is related to smoothness because technical noise shows up at very low frequencies, while we expect from semiconductors smooth $1/f$ noise in a wide frequency range. Smoothness was verified by comparison with a database of trusted spectra. Technical noise turned out to be a serious difficulty. As no data was found in the literature, we give some practical hints in Fig. 12.24.

The EOM requires a high microwave power (20 dBm or more), which is some 50 dB higher than the photodetector output. The isolation in the microwave circuits is hardly higher than about 120 dB. Thus crosstalk, influenced by the fluctuating dielectric constant of the environment, turns into a detectable signal. The system clearly senses the experimentalist waving a hand (≈ 0.2 m/s) at a distance of 3 m. The spectrum (Fig. 12.24.1, plot W) is easily taken for flicker. This problem can be mitigated using the new high-efficiency EOMs [?].

Air flow affects the delay of the optical fibers, thus some isolation is necessary to mitigate this effect. All our attempts failed until we inserted optical isolators in series with the photodetectors, and spliced all the fiber junctions (except the laser output). After this, the back-reflected light at the unused port of the coupler was below the sensitivity of the power-meter, which is 1 nW. Without isolation and splicing, individual spectra show spikes appearing at random times (Fig. 12.24.2, plot S). Averaging yields a smooth spectrum. Yet slope is incorrect (Fig. 12.24.3, plot A). Beside the mechanics of the connectors, we attribute this effect to reflection noise in the optical fibers [?, ?].

Even after isolating and splicing, we observed that bending a fiber may result in increased flickering. Afterwards, the spectrum may become irregular, or still be smooth with a clean $1/f$ slope, as in Fig. 12.24.4, plot F, but nevertheless incorrect. We interpret this as a change in the interference pattern in the fiber due to polarization. The observed increase in noise is clearly systematic, although reproducing the numerical value takes some effort.

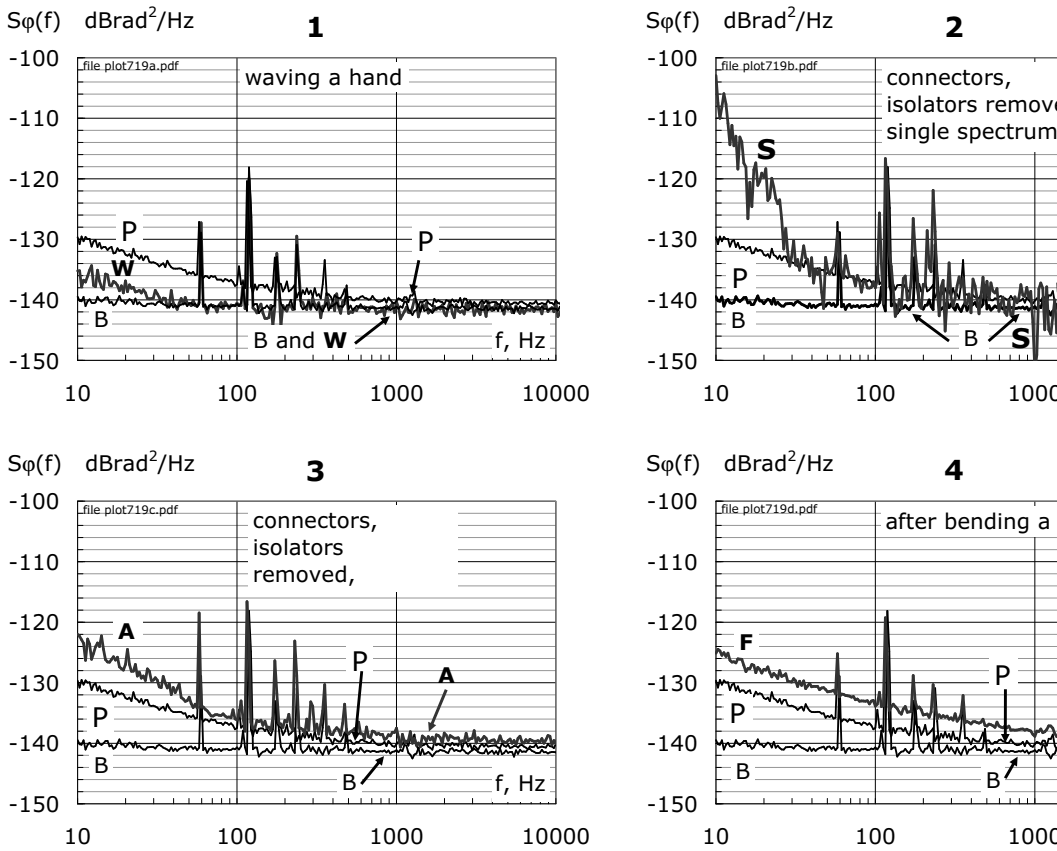


Figure 12.24: Examples of environment effects and experimental mistakes around the corner. For reference, all the plots show the instrument Background noise (spectrum B) and the noise spectrum of the Photodiode pair (spectrum P. Plot 1 spectrum W: the experimentalist Waves a hand gently (≈ 0.2 m/s), 3 m far away from the system. Plot 2 spectrum S: the optical isolators are removed and the connectors are restored at the input of the photodiodes (Single spectrum). Plot 3 spectrum A: same as plot 3, but Average spectrum. Plot 4 spectrum F: a Fiber is bended with a radius of ≈ 5 cm, which is twice that of a standard reel.

Spectral lines at 60 Hz and its multiples are present in the noise spectra, originated by magnetic fields, in all cases lower than $-110 \text{ dB}[\text{rad}^2]/\text{Hz}$. The level of these stray signals is about the same found routinely in the phase noise measurement with the saturated mixer method, yet with a carrier power of some 10 dBm instead of the -26 dBm of our experiments, thus with a signal-to-noise ratio proportionally higher. The superior immunity of the bridge scheme is due to microwave amplification of the noise sidebands before detecting.

The $1/f$ spectra of the detectors we measured are similar, and a value of $-120 \text{ dB}[\text{rad}^2]/\text{Hz}$ at $f = 1 \text{ Hz}$ can be taken as representative of both amplitude

and phase noise. Using the formulae available in [?, ?, ?], a spectrum of the form h_{-1}/f converted into the Allan (two-sample) variance $\sigma^2(\tau)$ is $\sigma^2 = 2 \ln(2) h_{-1}$ independent of the measurement time τ . The length of 1 rad R in a fiber of refraction index $n = 1.45$, at the modulation frequency $\nu_\mu = 9.9$ GHz, of is 3.3 mm. Thus a phase noise of -120 dBrad²/Hz at $f = 1$ Hz ($h_{-1} = 10^{-12}$) is equivalent to a fluctuation $\sigma_l(\tau) = 3.9$ nm of the optical length l .

12.18 Final remarks

It is generally accepted [Sikula, 2003] that flicker noise is an elusive phenomenon and that our understanding is based on models, the most accredited of which are due to Hooke [?] and to McWhorter [?], rather than on a unified theory. On the other hand, the presence of the phase and amplitude flickering in a microwave carrier is believed to be the dc flicker, up-converted by a nonlinearity. This also applies to the photodiode, even though in this case the dc bias exists only in the presence of light. In fact, removing the modulation results in a white microwave spectrum, flat around any frequency in the passband of the system.

The experimental difficulties we encountered are due to various forms of technical noise, at an exceedingly low level, which nevertheless may exceed the detector noise, unless great care is taken. On one hand, this means that the environment in which the diode is inserted must be revisited if one needs the lowest achievable noise. On the other hand, this means that the photodiode exhibits low noise and high stability, and that it has an unexploited potential for new and emerging applications.

Acknowledgements

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract of the National Aeronautics and Space Administration, and with support from ARL and AOSP/DARPA. We thank the Université Henri Poincaré for partially supporting E. Rubiola while visiting JPL, and F. Lardet-Vieudrin of FEMTO-ST, France, for providing low-flicker dc preamplifiers.

12.19 Flicker and fibers

Something to be done⁶

Jennifer A. Taylor+*, Frank Quinlan*, Archita Hati*, Craig Nelson*, Scott A. Diddams*, Shubhashish Datta, Abhay Joshi Phase Noise in the Photodetection of Ultra-short Optical Pulses

Dahyeon Lee, Takuma Nakamura, Jizhao Zang, Joe C. Campbell, Scott A. Diddams, Franklyn Quinlan, Reduction of Flicker Phase Noise in High-Speed Photode-

⁶Still to be done




Fig. 1 Time response of photo-diode HLPD. *Top*: typical time record for various optical energies per pulse (from *bottom* to *top*: 3.2, 6.4, 12.8, 19.2, 25.6, 32, 38.4, 44.8, 51.2, 57.6, and 64 pJ/pulse). *Bottom left* y-scale: peak voltage vs. optical energy per pulse (the corresponding *solid line* is a double exponential fit), *bottom right* y-scale: integrated electrical energy vs. optical energy per pulse (*the solid line* is a linear fit, exhibiting 78% quantum efficiency)

Figure 12.25: Zhang W, Li T, Lours M, Seidelin S, Santarelli G, Le Coq Y - Amplitude to phase conversion in InGaAs pin photodiodes etc - Appl Phys B 106, September 2011

tectors Under Ultrashort Pulse Illumination, IEEE Photonic J Volume 13, Number 3, June 2021.

How to make the propagation time through an optical fiber fully insensitive to temperature variations, ERIC NUMKAM FOKOUA,* MARCO N. PETROVICH, TOM BRADLEY, FRANCESCO POLETTI, DAVID J. RICHARDSON, AND RADAN

SLAVÍK, Optica 2017

Gianni Di Domenico,* Stéphane Schilt, and Pierre Thomann, Simple approach to the relation between laser frequency noise and laser line shape

Chapter 13

AM Noise

The close-in AM noise is often neglected, under the assumption that it is a minor problem as compared to phase noise. With the progress of technology and of experimental science, this assumption is no longer true. Yet, information in the literature is scarce or absent.

This report describes the measurement of the AM noise of rf/microwave sources in terms of $S_\alpha(f)$, i.e., the power spectrum density of the fractional amplitude fluctuation α . The proposed schemes make use of commercial power detectors based on Schottky and tunnel diodes, in single-channel and correlation configuration.

There follow the analysis of the front-end amplifier at the detector output, the analysis of the methods for the measurement of the power-detector noise, and a digression about the calibration procedures.

The measurement methods are extended to the relative intensity noise (RIN) of optical beams, and to the AM noise of the rf/microwave modulation in photonic systems.

Some rf/microwave synthesizers and oscillators have been measured, using correlation and moderate averaging. As an example, the flicker noise of a low-noise quartz oscillator (Wenzel 501-04623E) is $S_\alpha = 1.15 \times 10^{-13}/f$, which is equivalent to an Allan deviation of $\sigma_\alpha = 4 \times 10^{-7}$. The measurement systems described exhibit the world-record lowest background noise.

13.1 Basics

A quasi-perfect rf/microwave sinusoidal signal can be written as

$$v(t) = V_0 [1 + \alpha(t)] \cos[2\pi\nu_0 t + \phi(t)] \quad , \quad (13.1)$$

where $\alpha(t)$ is the fractional amplitude fluctuation, and $\phi(t)$ is the phase fluctuation. Equation (13.1) defines $\alpha(t)$ and $\phi(t)$. In low noise conditions, that is, $|\alpha(t)| \ll 1$ and $|\phi(t)| \ll 1$, Eq. (13.1) is equivalent to

$$\begin{aligned} v(t) &= V_0 \cos(2\pi\nu_0 t) + v_c(t) \cos(2\pi\nu_0 t) - v_s(t) \sin(2\pi\nu_0 t) \\ \text{with} \quad \alpha(t) &= \frac{1}{V_0} v_c(t) \quad \text{and} \quad \phi(t) = \frac{1}{V_0} v_s(t) . \end{aligned} \quad (13.2)$$

We make the following assumptions about $v(t)$, in agreement with actual cases of interest:

1. The expectation of the amplitude is V_0 . Thus $\mathbb{E}\{\alpha(t)\} = 0$.
2. The expectation of the frequency is ν_0 . Thus $\mathbb{E}\{\dot{\phi}(t)\} = 0$.
3. Low noise. $|\alpha(t)| \ll 1$ and $|\phi(t)| \ll 1$.
4. Narrow band. The bandwidth of α and ϕ is $B_\alpha \ll \nu_0$ and $B_\phi \ll \nu_0$.

It is often convenient to describe the close-in noise in terms of the single-side¹ power spectrum density $S(f)$, as a function of the Fourier frequency f . A model that has been found useful to describe $S(f)$ is the power-law $S(f) = \sum_i h_i f^i$. In the case of amplitude noise, generally the spectrum contains only the white noise $h_0 f^0$, the flicker noise $h_{-1} f^{-1}$, and the random walk $h_{-2} f^{-2}$. Accordingly,

$$S_\alpha(f) = h_0 + h_{-1} f^{-1} + h_{-2} f^{-2} \quad . \quad (13.3)$$

Random walk and higher-slope phenomena, like drift, are often induced by the environment. It is up to the experimentalist to judge the effect of environment.

The spectrum density can be converted into Allan variance using the formulae of Table 13.1.

The signal power is

$$P = \frac{V_0^2}{2R} (1 + \alpha)^2 \quad (13.4)$$

thus

$$P \simeq \frac{V_0^2}{2R} (1 + 2\alpha) \quad \text{because } \alpha \ll 1 \quad (13.5)$$

¹Most experimentalists prefer the single-side power spectrum density because all instruments work in this way. This is because the power can be calculated as $P = \int_0^B S(f) df$, which is far more straightforward than integrating over positive and (to some extent, mysterious) negative frequencies.

Table 13.1: Relationships between power spectrum density and Allan variance.

noise type	Spectrum density $S_\alpha(f)$	Allan variance $\sigma_\alpha^2(\tau)$
white	h_0	$\frac{h_0}{2\tau}$
flicker	$h_{-1}f^{-1}$	$h_{-1} \frac{2 \ln(2)}{4\pi^2}$
random walk	$h_{-2}f^{-2}$	$h_{-2} \frac{4\pi^2}{6} \tau$

It is convenient to rewrite P as $P = P_0 + \delta P$, with

$$P_0 = \frac{V_0^2}{2R} \quad \text{and} \quad \delta P \simeq 2P_0\alpha \quad (13.6)$$

The amplitude fluctuations are measured through the measurement of the power fluctuation δP ,

$$\alpha(t) = \frac{1}{2} \frac{\delta P}{P_0} \quad (13.7)$$

and of its power spectrum density,

$$S_\alpha(f) = \frac{1}{4} S_{\frac{\delta P}{P_0}}(f) = \frac{1}{4P_0^2} S_P(f) . \quad (13.8)$$

The measurement of a two-port device, like an amplifier, is made easy by the availability of the reference signal sent to the device input. In this case, the bridge (interferometric) method [?] enables the measurement of amplitude noise and phase noise with outstanding sensitivity. Yet, the bridge method can not be exploited for the measurement of the AM noise of oscillators, synthesizers and other signal sources. Other methods are needed, based on power detectors and on suitable signal processing techniques.

13.2 Single channel measurement

Figure 13.1 shows the basic scheme for the measurement of AM noise. The detector characteristics (Sec. 13.4) is $v_d = k_d P$, hence the ac component of the detected signal is $\tilde{v}_d = k_d \delta P$. The detected voltage is related to α by $\tilde{v}_d = k_d P_0 \frac{\delta P}{P_0}$, that is,

$$\tilde{v}_d(t) = 2k_d P_0 \alpha(t) . \quad (13.9)$$

Turning voltages into spectra, the above becomes

$$S_v(f) = 4k_d^2 P_0^2 S_\alpha(f) . \quad (13.10)$$

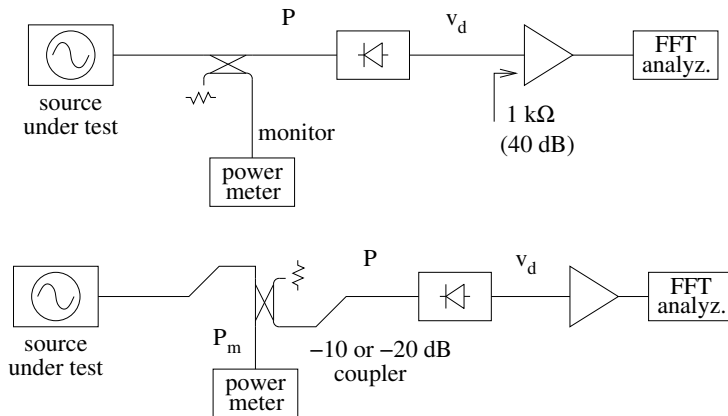


Figure 13.1: Top: Basic scheme for the measurement of AM noise. Bottom: A variant useful in some cases where the power detector operates at low power, while the power meter does not.

Therefore, the spectrum of α can be measured using

$$S_{\alpha}(f) = \frac{1}{4k_d^2 P_0^2} S_v(f) \quad . \quad (13.11)$$

Due to linearity of the network that precedes the detector (directional couplers, cables, etc.), the fractional power fluctuation $\delta P/P_0$ is the same in all the circuit, thus α is the same. As a consequence, the separate measurement of the oscillator power and of the attenuation from the oscillator to the detector is not necessary. The straightforward way to use Eq. (13.10), or (13.11), is to refer P_0 at detector input, and v_d at the detector output.

Interestingly, phase noise has virtually no effect on the measurement. This happens because the bandwidth of the detector is much larger than the maximum frequency of the Fourier analysis, hence no memory effect takes place.

In single-channel measurements, the background noise can only be assessed by measuring a low-noise source². Of course, this measurement gives the total noise of the source and of the instrument, which can not be divided. The additional hypothesis is therefore required, that the amplitude noise of the source is lower than the instrument background. Unfortunately, a trusted source will be hardly available in practice.

Calibration is needed, which consists of the measurement of the product $k_d P_0$. See Section 13.11.

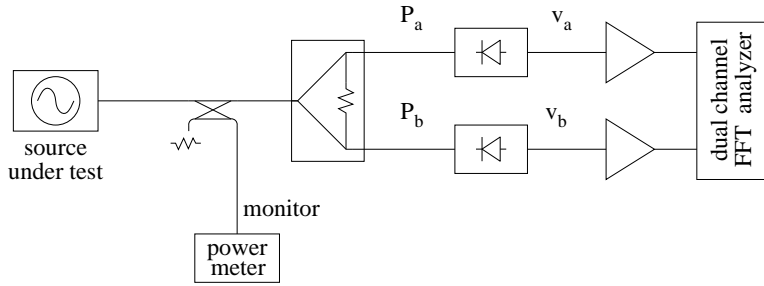


Figure 13.2: Correlation AM noise measurement.

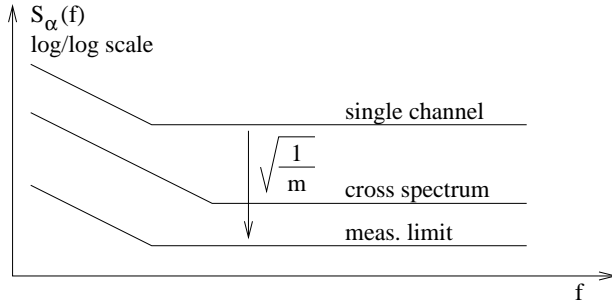


Figure 13.3: Spectra of the correlation AM noise measurement.

13.3 Dual channel (correlation) measurement

Figure 13.2 shows the scheme for the correlation measurement of AM noise. The signal is split into two branches, and measured by two separate power detectors and amplifiers. Under the assumption that the two channels are independent, the cross spectrum $S_{ba}(f)$ is proportional to $S_\alpha(f)$. In fact, the two dc signals are $v_a = k_a P_a \alpha$ and $v_b = k_b P_b \alpha$. The cross spectrum is

$$S_{ba}(f) = 4k_a k_b P_a P_b S_\alpha(f) \quad , \quad (13.12)$$

from which

$$S_\alpha(f) = \frac{1}{4k_a k_b P_a P_b} S_{ba}(f) \quad . \quad (13.13)$$

Averaging over m spectra, the noise of the individual channels is rejected by a factor $\sqrt{2m}$ (Fig. 13.3), for the sensitivity can be significantly increased. A further advantage of the correlation method is that the measurement of $S_\alpha(f)$ is validated by the simultaneous measurement of the instrument noise limit, that is, the single-channel noise divided by $\sqrt{2m}$. This solves one of the major problems of the single-channel measurement, i.e., the need of a trusted low-noise source.

²The reader familiar with phase noise measurements is used to measure the instrument noise by removing the device under test. This is not possible in the case of the AM noise of the oscillator.




Figure 13.4: Scheme of the diode power detector.

Larger is the power delivered by the source under test, larger is the instrument gain. This applies to single-channel measurements, where the gain is $4k_d^2 P_0^2$ [Eq. (13.10)], and to correlation measurements, where the gain is $4k_a k_b P_a P_b$ [Eq. (13.12)]. Yet in a correlation system the total power P_0 is split into the two channels, for $P_a P_b = \frac{1}{4} P_0^2$. Hence, switching from single-channel to correlation the gain drops by a factor $\frac{1}{4}$ (−6 dB). Let us now compare a correlation system to a single-channel system under the simplified hypothesis that the background noise referred at the detector output is unchanged. This happens if the noise of the dc preamplifier is dominant. In such cases, the background noise referred to the instrument input, thus to S_α , is multiplied by a factor $\frac{4}{\sqrt{2m}}$. The numerator “4” arises from the reduced gain, while the denominator $\sqrt{2m}$ is due to averaging. Accordingly, it must be $m > 8$ for the correlation scheme to be advantageous in terms of sensitivity. On the other hand, if the power of the source under test is large enough for the system to work at full gain in both cases, the dual-channel system exhibits higher sensitivity even at $m = 1$.

Calibration is about the same as for the single-channel measurements. See Section 13.11.

In laboratory practice, the availability of a dual-channel FFT analyzer is the most frequent critical point. If this instrument is available, the experimentalist will prefer the correlation scheme in virtually all cases.

13.4 Schottky and tunnel diode power detectors

A rf/microwave power detector uses the nonlinear response of a diode to turn the input power P into a dc voltage v_d . The transfer function is

$$v_d = k_d P \quad , \quad (13.14)$$

which defines the detector gain k_d . The physical dimension of k_d is A^{-1} . The technical unit often used in data sheets is mV/mW, equivalent to A^{-1} . The diodes can only work at low input level. Beyond a threshold power, the output voltage differs smoothly from Eq. (13.14). The actual response depends on the diode type.

Figure 13.4 shows the scheme of actual power detectors. The input resistor matches the high input impedance of the diode network to the standard value $R_0 = 50 \Omega$ over the bandwidth and over the power range. The value depends on the specific detector. The output capacitor filters the video³ signal, eliminating carrier from the output. A

³From the early time of electronics, the term ‘video’ is used (as opposed to ‘audio’) to emphasize the large bandwidth of the demodulated signal, regardless of the real purposes.

Table 13.2: some power-detector manufacturers (non-exhaustive list).

manufacturer	web site
Aeroflex/Metelics	aeroflex-metelics.com
Agilent Technologies	agilent.com
Advanced Control Components	advanced-control.com
Advanced Microwave	advancedmicrowaveinc.com
Eclipse	eclipsemicrowave.com
Herotek	herotek.com
Microphase	microphase.com/military/detectors.shtml
Omniiyg	omniyg.com
RLC Electronics	rlcelectronics.com/detectors.htm
S-Team	s-team.sk

Table 13.3: Typical characteristics of Schottky and tunnel power detectors.

parameter	Schottky	tunnel
input bandwidth	up to 4 decades 10 MHz to 20 GHz	1–3 octaves up to 40 GHz
VSVR max.	1.5:1	3.5:1
max. input power (spec.)	−15 dBm	−15 dBm
absolute max. input power	20 dBm or more	20 dBm
output resistance	1–10 kΩ	50–200 Ω
output capacitance	20–200 pF	10–50 pF
gain	300 V/W	1000 V/W
cryogenic temperature	no	yes
electrically fragile	no	yes

low capacitance makes the detector fast. On the other hand, a higher capacitance is needed if the detector is used to demodulate a low-frequency carrier. The two-diode configuration provides larger output voltage and some temperature compensation.

Power detectors are available off-the-shelf from numerous manufacturers, some of which are listed on Table 13.2. Agilent Technologies provides a series of useful application notes [Agilent-AN1449, 2003] about the measurement of rf/microwave power.

Two types of diode are used in practice, Schottky and tunnel. Their typical characteristics are shown in Table 13.3.

Schottky detectors are the most common ones. The relatively high output resistance and capacitance makes the detector suitable to low-frequency carriers, starting from some 10 MHz (typical). In this condition the current flowing through the diode is small, and the input matching to $R_0 = 50\ \Omega$ is provided by a low value resistor. Thus, the VSWR is close to 1:1 in a wide frequency range. Most of the input power




Figure 13.5: Response of a two-diode power detector.

is dissipated in the input resistance, which reduces the risk of damage in case of overload. A strong preference for negative output voltage seems to derive from the lower noise of P type Schottky diodes, as compared to N type ones, in conjunction with practical issues of mechanical layout. Figure 13.5 shows the response of a two-diode Schottky power detector. The quadratic response [Eq. (13.14)] derives from the diode resistance R_d , which is related to the saturation current I_0 by

$$R_d = \frac{\eta V_T}{I_0} , \quad (13.15)$$

where $\eta \in [1 \dots 2]$ is a parameter that derives from the junction technology; $V_T = kT/q \simeq 25.6$ mV at room temperature is the thermal voltage. At higher input level, R_d becomes too small and the detector response turns smoothly from quadratic to linear, like the response of the common AM demodulators and power rectifiers.

Tunnel detectors are actually *backward* detectors. The backward diode is a tunnel diode in which the negative resistance in the forward-bias region is made negligible by appropriate doping, and used in the reverse-bias region. Most of the work on such detectors dates back to the sixties [?, ?, Hall, 1960]. Tunnel detectors exhibit fast switching and higher gain than the Schottky counterpart. A low output resistance is necessary, which affects the input impedance. Input impedance matching is therefore poor. In the measurement of AM noise, as in other applications in which fast response is not relevant, the output resistance can be higher than the recommended value, and limited only by noise considerations. At higher output resistance the gain further increases. Tunnel diodes also work in cryogenic environment, provided the package tolerates the mechanical stress of the thermal contraction.

Figures 13.6–13.7 and Table 13.4 show the conversion gain of two detectors, measured at the FEMTO-ST Institute. As expected, the Schottky detector leaves smoothly the quadratic law (true power detection) at some -12 dBm, where it becomes a peak voltage detector. The response of the tunnel detector is quadratic up




Figure 13.6: Measured response of a Schottky detector Herotek DZR124AA (Serial X-110).

Table 13.4: Measured conversion gain.

load resistance, Ω	detector gain, A^{-1}	
	DZR124AA (Schottky)	DT8012 (tunnel)
	s. no. 227489	s. no. 232028)
1×10^2	35	292
3.2×10^2	98	505
1×10^3	217	652
3.2×10^3	374	724
1×10^4	494	750
conditions: power -50 to -20 dBm		

to a maximum power lower than that of the Schottky diode. This is due to the lower threshold of the tunnel effect. The output voltage shows a maximum at some 0 dBm, then decreases. This is ascribed to the tunnel-diode conduction in the forward region.

At the FEMTO-ST Institute, I routinely use the Herotek DZR124AA (Schottky) and DT8012 (tunnel). At the JPL, I have sometimes used the pair HP432A, and more recently the same Herotek types that I use at the LPMO. The old HP432A pair




Figure 13.7: Measured response of a tunnel detector Herotek DT8012 (Serial X-111).

available shows an asymmetry of almost 10 dB in flickering, which causes experimental difficulties. This asymmetry might be a defect of the specific sample.

13.5 The double balanced mixer

Some articles from the NIST [?] ⁴ suggest the use of a double balanced mixer as the power detector in AM noise measurements, with a configuration similar to that of Fig. 13.8. Higher sensitivity is obtained with cross-spectrum measurements, using two mixers. The double balanced mixer is operated in the following conditions:

1. Both RF and LO inputs are *not* saturated.
2. The RF and LO signals are in-phase.

At low frequencies the in-phase condition may be hard wired, omitting the variable phase shifter.

However useful, the available articles do not explain the diode operation in terms of electrical parameters. One can expect that the background noise of the mixer used as the AM detector can not be lower than that of a Schottky-diode power detector for the simple reason that the mixer contains a ring of Schottky diodes. One could guess that the use of the mixer for AM noise measurements originates from having

⁴Fred Walls published some articles in the FCS Proceedings. Craig Nelson showed the method in his tutorial given at the 2003 FCS.




Figure 13.8: A double balanced mixer can be used as the power detector for the measurement of AM noise. Data are from [?].

had mixers and phase shifter on hand for a long time at the NIST department of phase noise.

13.6 Power detector noise

Two fundamental types of noise are present in a power detector, shot noise and thermal noise [?, Sec. V]. In addition, detectors show flicker noise. The latter is not explained theoretically, for the detector characterization relies on experimental parameters. Some useful pieces of information are available in [?].

Owing to the shot effect, the average current \bar{i} flowing in the diode junction is affected by a fluctuation of power spectral density

$$S_i = 2q\bar{i} \quad \text{A}^2/\text{Hz}, \quad (13.16)$$

Using the Ohm law $v = Ri$ across the load resistor R , the noise voltage at the detector output is

$$S_v = 2qR\bar{v} \quad \text{V}^2/\text{Hz}. \quad (13.17)$$

Then, the shot noise is referred to the input-power noise using $v = k_d P$. Thus, at the




Figure 13.9: Power spectral density of the detector noise, referred at the input.

operating power P_0 it holds that

$$S_P = 2q \frac{RP_0}{k_d} \quad \text{shot noise, W}^2/\text{Hz}. \quad (13.18)$$

The thermal noise across load resistance R has the power spectral density

$$S_v = 4kTR \quad \text{V}^2/\text{Hz}, \quad (13.19)$$

which turns into

$$S_P = \frac{4kTR}{k_d^2} \quad \text{thermal noise, W}^2/\text{Hz}. \quad (13.20)$$

referred to the detector input. An additional thermal-noise contribution comes from the dissipative resistance of the diodes. This can be accounted for by increasing the value of R in Equations (13.19) and (13.20). It should be remarked that diode differential resistance is not a dissipative phenomenon, for there is no thermal noise associated to it.

Figure 13.9 shows the equivalent input noise as a function of power. The shot noise is equal to the thermal noise, $(S_P)_{\text{shot}} = (S_P)_{\text{thermal}}$, at the critical power

$$P_c = \frac{2}{k_d} \frac{kT}{q}. \quad (13.21)$$

It turns out that the power detector is always used in the low power region ($P_0 < P_c$), where shot noise is negligible. In fact, taking the data of Table 13.4 as typical of actual detectors, P_c spans from +2 dBm to -10 dBm for the Schottky diodes ($35 \text{ A}^{-1} < k_d < 500 \text{ A}^{-1}$), and from -8 dBm to -12 dBm for the tunnel diodes ($330 \text{ A}^{-1} < k_d < 820 \text{ A}^{-1}$), depending on the load resistance. On the other hand, the detector turns from the quadratic (power) response to the linear (voltage) response at a significantly lower power. This can be seen on Figure 13.6 and 13.7.

Looking at the specifications of commercial power detectors, information about noise is scarce. Some manufacturers give the NEP (Noise Equivalent Power) parameter, i.e., the power at the detector input that produces a video output equal to that of

the device noise. In no case is said whether the NEP increases or not in the presence of a strong input signal, which is related to precision. Even worse, no data about flickering is found in the literature or in the data sheets. Only one manufacturer (Herotek) claims the low flicker feature of its tunnel diodes, yet without providing any data.

The power detector is always connected to some kind of amplifier, which is noisy. Denoting with $(h_0)_{\text{ampli}}$ and $(h_{-1})_{\text{ampli}}$ the white and flicker noise coefficients of the amplifier, the spectrum density referred at the input is

$$S_P(f) = \frac{(h_0)_{\text{ampli}}}{k_d^2} + \frac{(h_{-1})_{\text{ampli}}}{k_d^d} \frac{1}{f}. \quad (13.22)$$

The amplifier noise coefficient $(h_0)_{\text{ampli}}$ is connected to the noise figure by $(h_0)_{\text{ampli}} = (F - 1)kT$. Yet we prefer not to use the noise figure because in general the amplifier noise results from voltage noise and current noise, which depends on R . Equation (13.22) is rewritten in terms of amplitude noise using $\alpha = \frac{1}{2} \frac{\delta P}{P_0}$ [Eq. (13.7)]. Thus,

$$S_\alpha(f) = \frac{1}{2P_0} \frac{qR}{k_d} + \frac{1}{P_0^2} \frac{kTR}{k_d^2} + \frac{1}{4P_0^2} \frac{(h_0)_{\text{ampli}}}{k_d^2} + \frac{1}{4P_0^2} \frac{(h_{-1})_{\text{ampli}}}{k_d^2} \frac{1}{f}. \quad (13.23)$$

After the first term of Eq. (13.22), the critical power becomes

$$P'_c = \frac{2}{k_d} \frac{kT}{q} + \frac{(h_0)_{\text{ampli}}}{2qRk_d}. \quad (13.24)$$

This reinforces the conclusion that in actual conditions the shot noise is negligible.

13.7 Design of the front-end amplifier

For optimum design, one should account for the detector noise and for the noise of the amplifier, and find the most appropriate amplifier and operating conditions. Yet, the optimum design relies upon the detailed knowledge of the power-detector noise, which is one of our targets (Sec. 13.8). Thus, we provisionally neglect the excess noise of the power detector. The first design is based on the available data, i.e., thermal noise and the noise of the amplifier. Operational amplifiers or other types of impedance-mismatched amplifiers are often used in practice. As a consequence, a single parameter, i.e., the noise figure or the noise temperature, is not sufficient to describe the amplifier noise. Voltage and current fluctuations must be treated separately, according to the popular Rothe-Dahlke model [?] (Fig. 13.10). The amplifier noise contains white and flicker, thus

$$(S_v)_{\text{ampli}} = h_{0,v} + h_{-1,v} \frac{1}{f} \quad (13.25)$$

$$(S_i)_{\text{ampli}} = h_{0,i} + h_{-1,i} \frac{1}{f}. \quad (13.26)$$

The design can be corrected afterwards, accounting for the flicker noise of the detector.




Figure 13.10: Rothe-Dahlke model of the noisy amplifier.




Figure 13.11: The load resistor turns the current noise into fully-correlated noise.

Single-channel systems

Accounting for shot and thermal noise, and for the noise of the amplifier, the noise spectrum density is

$$S_v = 2qR\bar{v} + 4kTR + (S_v)_{\text{ampli}} + R^2 (S_i)_{\text{ampli}} \quad (13.27)$$

at the amplifier input, and

$$S_P = 2q \frac{RP}{k_d} + \frac{4kTR}{k_d^2} + \frac{(S_v)_{\text{ampli}}}{k_d^2} + \frac{R^2 (S_i)_{\text{ampli}}}{k_d^2} \quad (13.28)$$

referred to the rf input. The detector gain k_d depends on R , thus the residual S_P can not be arbitrarily reduced by decreasing R . Instead, there is an optimum R at which the system noise is at its minimum.

Correlation-and-averaging systems

The noise contribution of the amplifier can be reduced by measuring the cross spectrum at the output of two amplifiers connected to the power detector, provided that the noise of the amplifiers is independent. For this to be true, the optimum design of the front-end amplifier changes radically. Figure 13.11 points out the problem. The current noise of each amplifier turns into a random voltage fluctuation across the load

resistance R . Focusing only on the amplifier noise, the voltage at the two outputs is

$$\begin{aligned}v_a &= g(v_1 + Ri_1 + Ri_2) \\v_b &= g(v_2 + Ri_1 + Ri_2)\end{aligned}$$

The terms gv_1 and gv_2 are independent, for their contribution to the cross spectrum density is reduced by a factor $\frac{1}{\sqrt{2m}}$, where m is the number of averaged spectra. Conversely, a term

$$gR(i_1 + i_2)$$

is present at the two outputs. This term is exactly the same, thus it can not be reduced by correlation and averaging. Consequently, the lowest current noise is the most important parameter, even if this is obtained at expense of a larger voltage noise. Yet, the rejection of larger voltage noise requires large m , for some tradeoff may be necessary.

Examples

This section shows some design attempts, aimed at the lowest white and flicker noise at low Fourier frequencies, up to 0.1–1 MHz, where operational amplifiers can be exploited in a simple way.

A preliminary analysis reveals that, at the low resistance values required by the detector, BJT amplifiers perform lower noise than field-effect transistors. On the other hand, the noise rejection by correlation and average requires low current noise, for JFET amplifiers are the best choice. In fact, BJTs can not be used because of the current noise, while MOSFETs show $1/f$ noise significantly larger (10 dB or more) than JFETs.

Using two detectors (DZR124AA and DT8012), we try the operational amplifiers and transistor pairs listed on Table 13.5. These amplifiers are selected with the criterion that each one is a low-noise choice in its category.

AD743 and OPA627 are general-purpose precision JFET amplifiers, which exhibit low bias current, hence low current noise. They are intended for correlation-and-averaging schemes (Fig. 13.11). The **OP625** is similar to the OP627 but for the frequency compensation, which enables unity-gain operations, yet at expenses of speed. It is used successfully in the measurement of the excess noise of semiconductors, where large averaging size is necessary in order to rid of the amplifier noise [?].

LT1028 is a fast BJT amplifier with high bias current in the differential input stage. This feature makes it suitable to low-noise applications in which the source resistance is low. In fact, the optimum noise resistance $R_b = \sqrt{h_v/h_i}$ is of 900 Ω for white noise, and of 105 Ω for flicker. These values are in the preferred range for proper operation of the power detectors.

MAT02 and MAT03 are bipolar matched pairs. They exhibit lower noise than operational amplifiers, and they are suitable to the design for low resistance

Table 13.5: Noise parameters of some selected amplifiers.

type	voltage		current		notes
	white $\sqrt{h_{0,v}}$	flicker $\sqrt{h_{-1,v}}$	white $\sqrt{h_{0,i}}$	flicker $\sqrt{h_{-1,i}}$	
AD743	2.9	18	6.9	0.1	jfet op-amp
LT1028	0.9	1.7	1000	16	bjt op-amp
MAT02	0.9	1.6	900	1.6	npn bjt matched pair
MAT03	0.7	1.2	1200	11	pnp bjt matched pair
OP27	3.0	4.3	400	4.7	bjt op-amp
OP177	10	8.0	125	1.6	bjt op-amp
OP1177	8.0	8.3	200	1.5	bjt op-amp
OPA627	4.5	45	2.5	—	jfet op-amp
unit	$\frac{\text{nV}}{\sqrt{\text{Hz}}}$	$\frac{\text{nV}}{\sqrt{\text{Hz}}}$	$\frac{\text{fA}}{\sqrt{\text{Hz}}}$	$\frac{\text{pA}}{\sqrt{\text{Hz}}}$	

of the source, like the LT1028. The MAT03 was successfully employed in the design of a low-noise amplifier optimized for $50\ \Omega$ sources [?].

OP27 and OP37 are popular general-purpose precision BJT amplifiers, most often used in low-noise applications. Their noise characteristics are about identical. The OP27 is fully compensated, for it is stable at closed-loop gain of one. The OP37 is only partially compensated, which requires a minimum closed-loop gain of five for stable operation. Of course, lower compensation increases bandwidth and speed.

OP177 and OP1177 are general-purpose precision BJT amplifiers with low bias current in the differential input stage, thus they exhibit lower current noise than other BJT amplifiers. They can be an alternative if the design based on JFET amplifier fails.

In the try-and-error process, we take into account shot noise, thermal noise, amplifier voltage noise, and amplifier current noise. The total white noise is the sum of all them

$$S_P = 2q \frac{RP}{k_d} + \frac{4kTR}{k_d^2} + \frac{h_{0,v}}{k_d^2} + \frac{R^2 h_{0,i}}{k_d^2} \quad \text{total white } [+]. \quad (13.29)$$

The total flicker accounts for the voltage and current of the amplifier

$$S_P = \frac{h_{-1,v}}{k_d^2} + \frac{R^2 h_{-1,i}}{k_d^2} \quad \text{total flicker } [\diamond]. \quad (13.30)$$

The correlated noise spectrum is equal to the total noise minus the voltage noise spectrum of the amplifier, which is independent

$$S_P = 2q \frac{RP}{k_d} + \frac{4kTR}{k_d^2} + \frac{R^2 h_{0,i}}{k_d^2} \quad \text{correlated white } [\bigcirc], \quad (13.31)$$

$$S_P = \frac{R^2 h_{-1,i}}{k_d^2} \quad \text{correlated flicker } [\square]. \quad (13.32)$$

The flicker noise of the detector, still not available, is to be added to Equations (13.30) and (13.32).

After evaluation, we discard the OP177 because the OP1177 exhibits superior performances in all noise parameters. Similarly, we discard the OP627 in favour of the AD743. The lower current noise of the OP627 can be exploited only at very large number of averages, which is impractical. On the other hand, the lower voltage noise of the AD743 helps to keep the number of averages reasonable. Figures 13.12 and 13.13 show a summary of the expected noise as a function of the load resistance R . The symbols $+$, \bigcirc , \diamond , \square have the same meaning in Equations (13.29)–(13.32) and in Figures 13.12 and 13.13. Analyzing the noise plots we restrict our attention to the detector-amplifier pairs of Table 13.6, in which we identify the following interesting configurations.

Detector DZR124AA (Schottky)

AD743 and 3.2 k Ω . Best design for a correlation-and-averaging system. In principle, lower white noise can be achieved with lower load resistance, yet only with $m > 10^4$.

MAT02 and 1 k Ω . Lowest flicker in real-time measurements (single-channel).

OP27 and 1 k Ω . Simple nearly-optimum design if one is interested to white noise. In fact, the white noise is 1.5 dB higher than that of the two above configurations; this gap increases only at $m > 300$ if the AD743 (jfet) is used. On the other hand, the flicker noise is high and can not reduced by correlation.

Detector DT8012 (tunnel)

AD743 and 1 k Ω . Best design for a correlation-and-averaging system. Slightly lower white noise can be obtained at lower R , yet at expenses of larger m and of larger flicker noise.

LT1028 and 100 Ω . Simple nearly-optimum design for real-time (single channel) systems. This configuration, as compared to the best one (MAT02 with 320 Ω load) shows white noise 1.2 dB higher, and flicker noise 9 dB higher.

MAT02 and 320 Ω . Lowest white and flicker noise in real-time measurements (single-channel).

MAT02 and 100 Ω . Close to the lowest white and flicker noise in real-time measurements (single-channel). Fairly good for correlation at moderate averaging, up to $m = 360$ for white noise, and $m = 90$ for flicker.

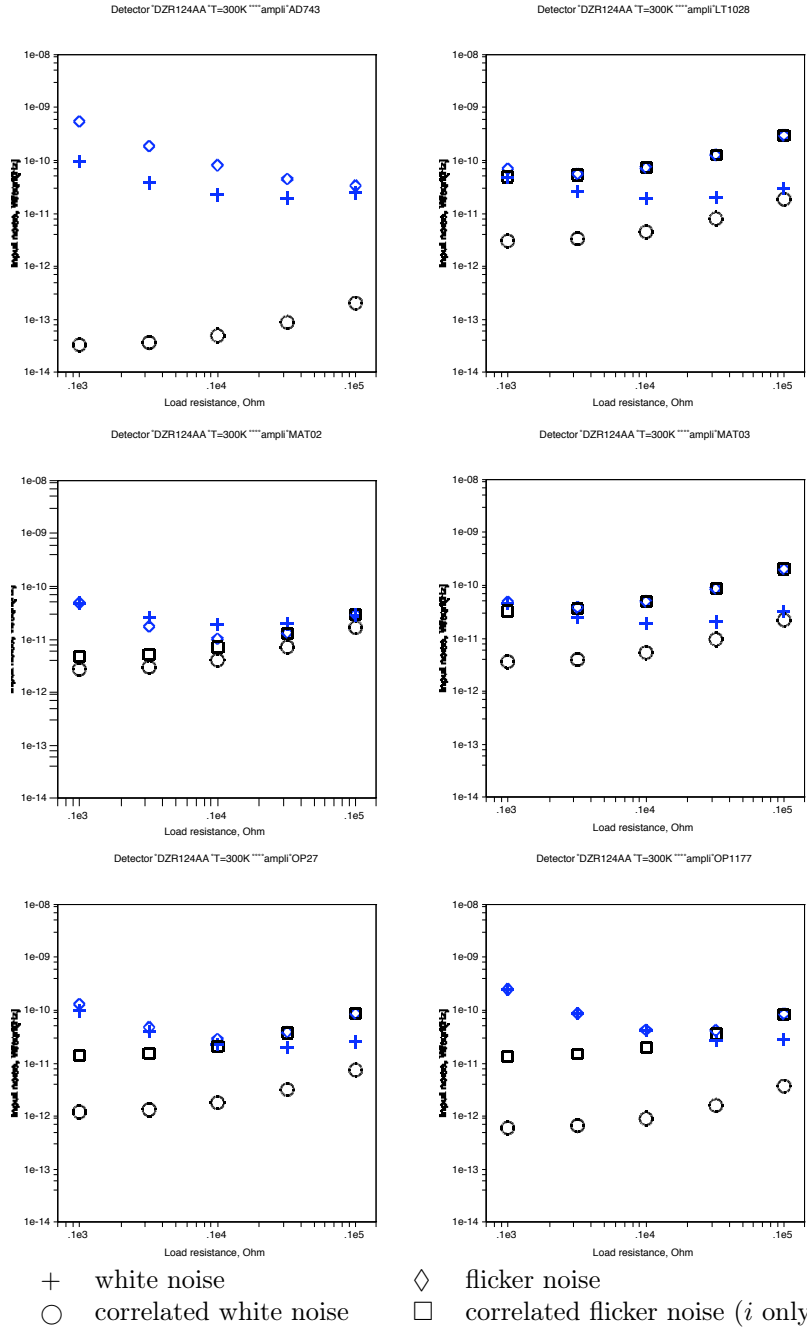


Figure 13.12: Expected noise of the DZR124AA (Schottky) followed an amplifier.

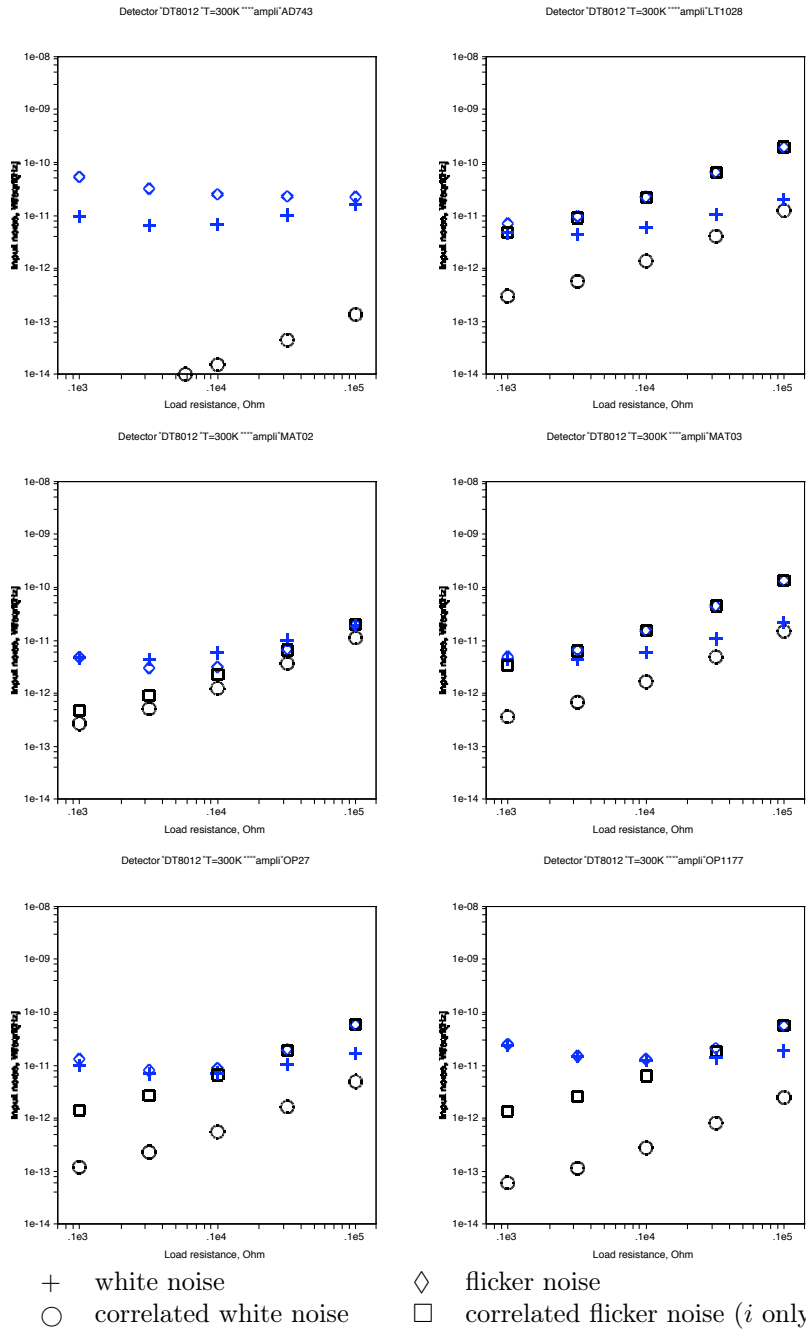


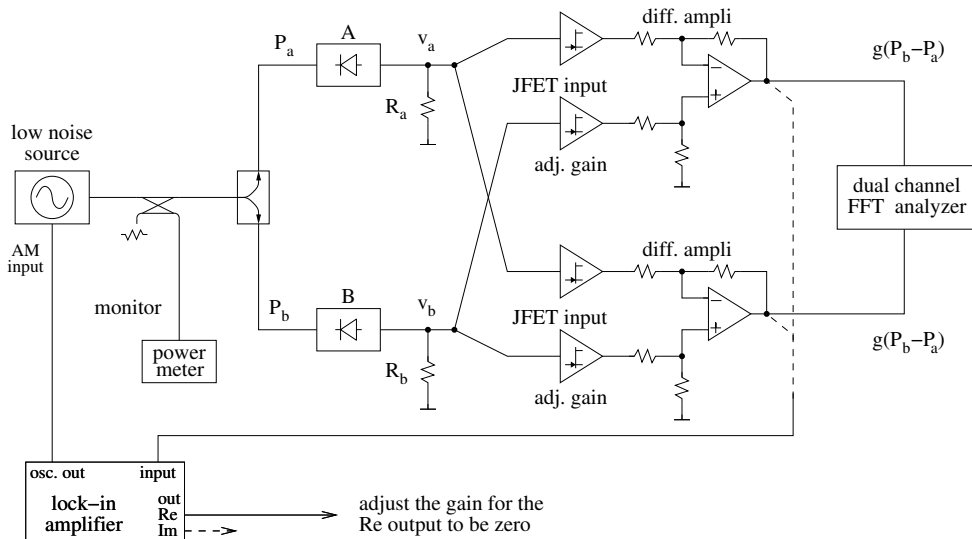
Figure 13.13: Expected noise of the DT8012 (Tunnel) followed an amplifier.

Table 13.6: Expected noise of a power detector followed by an amplifier. The grey box shows the best choice for the amplifier-detector pair.

Equivalent input noise							
detector and amplifier	noise type	load resistance					unit
		100	320	1000	3200	10000	Ω
DZR124AA and AD743	white	99.2	39.5	22.8	19.6	24.5	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	541	186	80.8	44.8	33.3	
	correl. white	0.050	0.042	0.052	0.089	0.205	
	correl. flicker	—	—	—	—	—	
DZR124AA and MAT02	white	54.7	27.5	19.6	19.8	29.2	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	48.6	17.7	10.4	13.5	29.9	
	correl. white	4.06	3.45	4.24	7.28	16.7	
	correl. flicker	7.22	6.13	7.54	12.9	29.8	
DZR124AA and OP27	white	102	40.3	23.1	20.0	25.6	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	131	48.0	29.4	39.5	87.8	
	correl. white	1.80	1.53	1.88	3.23	7.44	
	correl. flicker	21.2	18.0	22.1	38.0	87.4	
DT8012 and AD743	white	9.86	6.77	7.00	10.0	16.2	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	53.7	32.0	24.8	22.8	22.1	
	correl. white	0.005	0.007	0.016	0.045	0.135	
	correl. flicker	—	—	—	—	—	
DT8012 and LT1028	white	5.44	4.72	6.06	10.2	20.1	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	8.78	10.9	23.3	66.0	197	
	correl. white	0.448	0.657	1.45	4.12	12.3	
	correl. flicker	7.17	10.5	23.2	66.0	197	
DT8012 and MAT02	white	5.43	4.72	6.02	10.1	19.3	10^{-12} $\frac{\text{W}}{\sqrt{\text{Hz}}}$
	flicker	4.83	3.03	3.20	6.90	19.8	
	correl. white	0.403	0.591	1.30	3.71	11.1	
	correl. flicker	0.717	1.05	2.32	6.60	19.7	
White and shot noise, plus white and flicker noise of the amplifier. The flicker noise of the power detector is not accounted for.							

Remark. Generally, tunnel detectors show higher gain than Schottky detectors. The fact that they exhibit lower noise is a consequence. On the other hand, the Schottky detectors are often preferred because of wider bandwidth, and because of higher tolerance to electrical stress and to experimental errors.

A – Two-detector correlation measurement



B – Three-detector correlation measurement

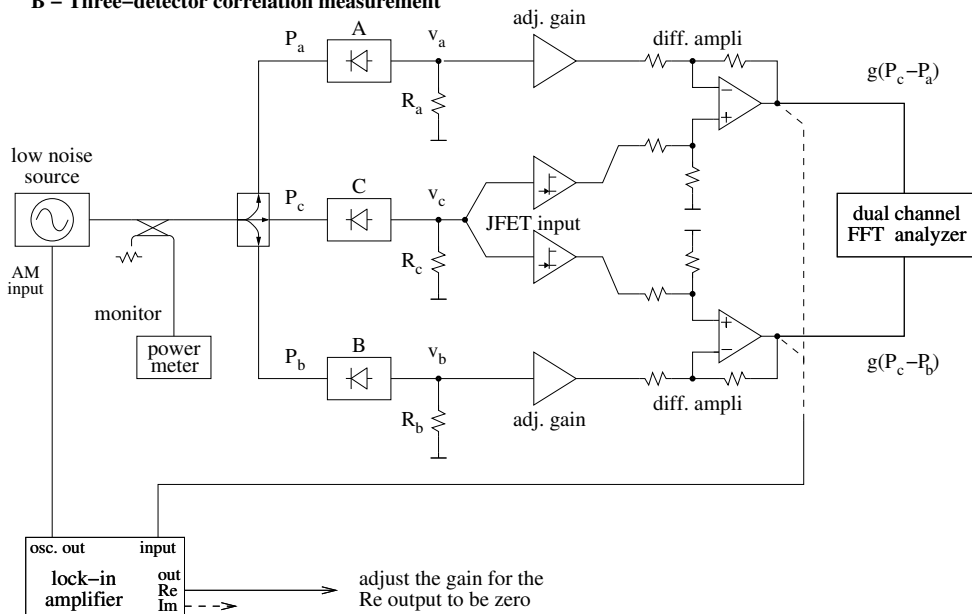


Figure 13.14: Measurement of the power detector noise.

13.8 The measurement of the power detector noise

A detector alone can be measured only if a reference source is available whose AM noise is lower than the detector noise, and if the amplifier noise can be made negligible.

These are unrealistic requirements.

It is useful to compare two detectors, as in Fig. 13.14 A. The trick is to measure a differential signal $g(P_b - P_a) \approx 0$, which is not affected by the power fluctuation of the source. The lock-in helps in making the output independent of the power fluctuations. Some residual PM noise has no effect on the detected voltage.

One problem with the scheme of Fig. 13.14 A is that the measured noise is the sum of the noise of the two detectors, for the result relies upon the assumption that the two detectors are about equal. This is fixed with the correlation scheme of Fig. 13.14 B). The detector c is the device under test, while the two other detectors are used to cancel the fluctuations of the input power. Thus

$$g(P_c - P_a) \quad \text{and} \quad g(P_c - P_b)$$

After rejecting the single-channel noise by correlation and averaging, there results the noise of the detector C.

Another problem with the schemes of Fig. 13.14 is that the noise of the amplifier is taken in. This is fixed by using two independent amplifiers at the output of the power detector, as in Fig. 13.15. In order to reject the current noise, these amplifiers must be of the JFET type. In the three detector scheme of Fig. 13.14 B, it is convenient to use BJT amplifiers in the reference branches (A and B), and a JFET amplifier in the branch C. The reason for this choice is the lower noise of the BJTs, which improves the measurement speed by reducing the minimum m needed for a given sensitivity.

13.9 AM noise in optical systems

Equation (13.1) also describes a quasi-perfect optical signal, under the same hypotheses 1–4 of page 206. The voltage $v(t)$ is replaced with the electric field. Yet, the preferred physical quantity used to describe the AM noise is the Relative Intensity Noise (RIN), defined as

$$\text{RIN} = S_{\frac{\delta I}{I_0}}(f) \quad , \quad (13.33)$$

that is, the power spectrum density of the normalized intensity fluctuation

$$\frac{(\delta I)(t)}{I_0} = \frac{I(t) - I_0}{I_0} \quad . \quad (13.34)$$

The RIN includes both fluctuation of power and the fluctuation of the power cross-section distribution. If the cross-section distribution is constant in time, the optical intensity is proportional to power

$$\frac{\delta I}{I_0} = \frac{\delta P}{P_0} \quad . \quad (13.35)$$

In optical-fiber system, where the detector collects all the beam power, the term RIN is improperly used for the relative power fluctuation. Reference [?] analyzes on the origin of RIN in semiconductor lasers, while References [?, ?] provide information on some topics of measurement.

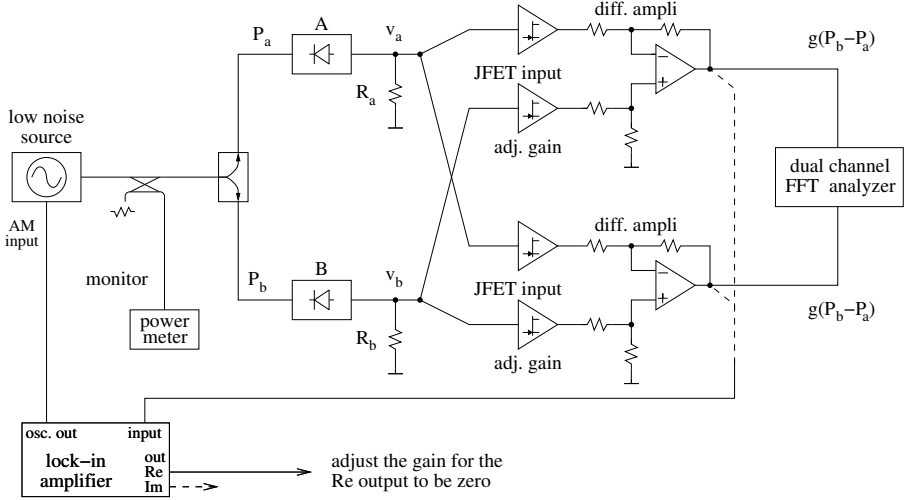
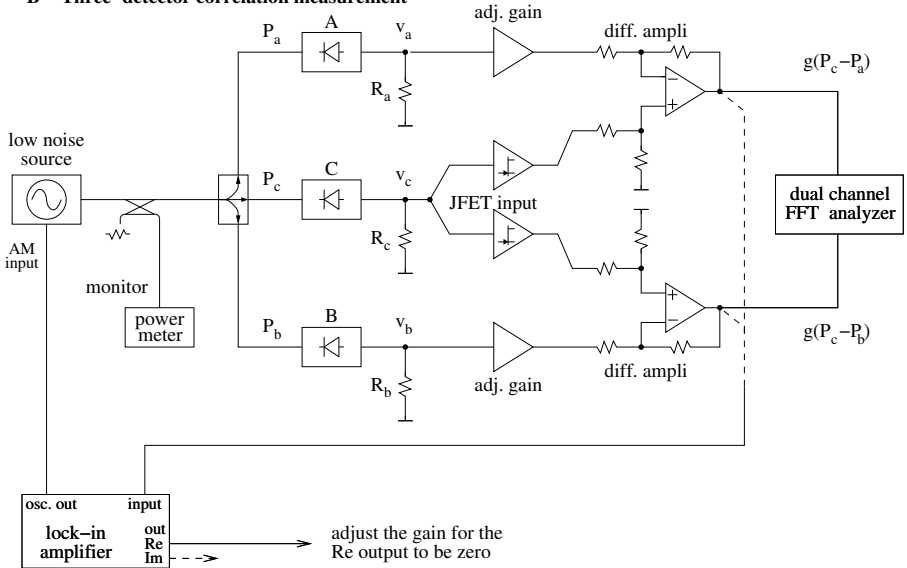
A – Two-detector correlation measurement**B – Three-detector correlation measurement**

Figure 13.15: Improved scheme for the measurement of the power detector noise.

In low-noise conditions, $|\delta I/I_0| \ll 1$, and assuming that the cross-section distribution is constant, the power fluctuations are related to the fractional amplitude noise α by

$$\frac{\delta I}{I_0} = \frac{\delta P}{P_0} = 2\alpha \quad , \quad (13.36)$$

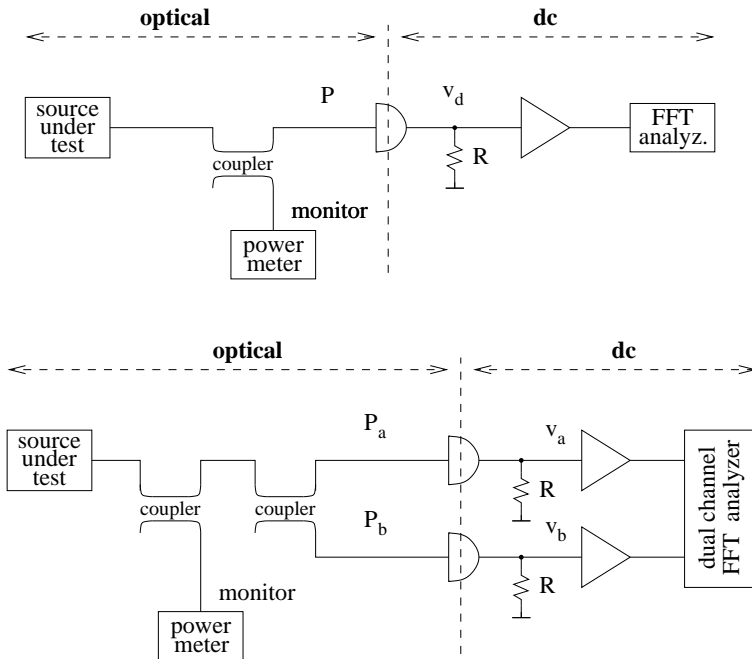


Figure 13.16: RIN measurement in optical-fiber systems. In a traditional system, beam splitters are used instead of the couplers.

thus

$$\text{RIN}(f) = 4S_{\alpha}(f) \quad . \quad (13.37)$$

Generally laser sources show a noise spectrum of the form

$$\text{RIN}(f) = h_0 + h_{-1}f^{-1} + h_{-2}f^{-2} \quad , \quad (13.38)$$

in which the flicker noise can be hidden by the random walk. Additional fluctuations induced by the environment may be present.

Figure 13.16 shows two measurement schemes. The output signal of the photodetector is a current proportional to the photon flux. Accordingly, the gain parameter is the detector responsivity ρ , defined by

$$i = \rho P \quad , \quad (13.39)$$

where

$$i = \frac{q\eta P}{h\nu} \quad (13.40)$$

is the photocurrent, thus

$$\rho = \frac{q\eta}{h\nu} \quad . \quad (13.41)$$

The schemes of Fig. 13.16 are similar to Fig. 13.1 (single-channel) and Fig. 13.2 (cross-spectrum). The dual-channel scheme is preferred because of the higher sensitivity, and because it makes possible to validate the measurement through the number of averaged spectra.

Noise is easily analyzed with the methods shown in Section 13.7. Yet in this case the virtual-ground amplifier is often preferred, which differs slightly from the examples shown in Section 13.7. A book [?] is entirely devoted to the special case of the photodiode amplifier.

13.10 AM noise in microwave photonic systems

Microwave and rf photonics is being progressively recognized as an emerging domain of technology [?, ?]. It is therefore natural to investigate in noise in these systems.

The power⁵ $P_\lambda(t)$ of the optical signal is sinusoidally modulated in intensity at the microwave frequency ν_μ is

$$P_\lambda(t) = \bar{P}_\lambda (1 + m \cos 2\pi\nu_\mu t) , \quad (13.42)$$

where m is the modulation index⁶. Eq. (13.42) is similar to the traditional AM of radio broadcasting, but optical power is modulated instead of RF voltage. In the presence of a distorted (nonlinear) modulation, we take the fundamental microwave frequency ν_0 . The detector photocurrent is

$$i(t) = \frac{q\eta}{h\nu_\lambda} \bar{P}_\lambda (1 + m \cos 2\pi\nu_\mu t) , \quad (13.43)$$

where η the quantum efficiency of the photodetector. The oscillation term $m \cos 2\pi\nu_\mu t$ of Eq. (13.43) contributes to the microwave signal, the term “1” does not. The microwave power fed into the load resistance R_0 is $\bar{P}_\mu = R_0 \bar{i}^2$, hence

$$\bar{P}_\mu = \frac{1}{2} m^2 R_0 \left(\frac{q\eta}{h\nu_\lambda} \bar{P}_\lambda \right)^2 . \quad (13.44)$$

The discrete nature of photons leads to the shot noise of power spectral density $2qiR$ [W/Hz] at the detector output. By virtue of Eq. (13.43),

$$N_s = 2 \frac{q^2 \eta}{h\nu_\lambda} \bar{P}_\lambda R \quad (\text{shot noise}) . \quad (13.45)$$

In addition, there is the equivalent input noise of the amplifier loaded by R , whose power spectrum is

$$N_t = FkT \quad (\text{thermal noise and amplifier noise}) , \quad (13.46)$$

⁵In this section we use the subscript λ for ‘light’ and μ for ‘microwave’.

⁶We use the symbol m for the modulation index, as in the general literature. There is no ambiguity because the number of averages (m) is not used in this section.

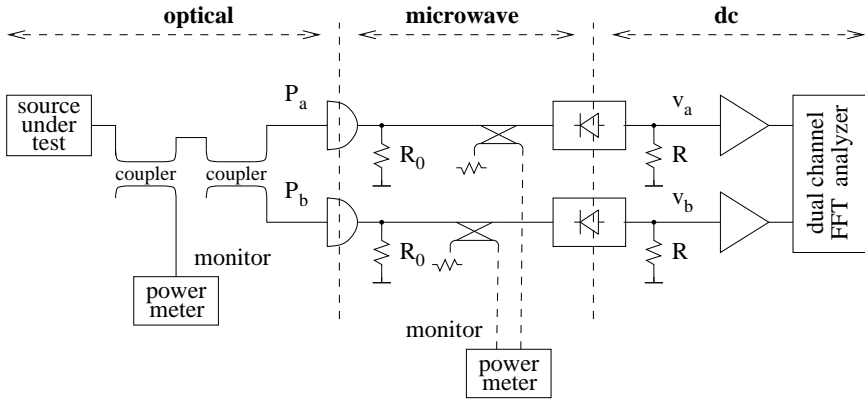


Figure 13.17: Measurement of the microwave AM noise of a modulated light beam.

where F is the noise figure of the amplifier, if any, at the output of the photodetector. The white noise $N_s + N_t$ turns into a noise floor

$$S_\alpha = \frac{N_s + N_t}{\bar{P}_\mu} . \quad (13.47)$$

Using (13.44), (13.45) and (13.46), the floor is

$$S_\alpha = \frac{2}{m^2} \left[2 \frac{h\nu_\lambda}{\eta} \frac{1}{\bar{P}_\lambda} + \frac{FkT}{R} \left(\frac{h\nu_\lambda}{q\eta} \right)^2 \left(\frac{1}{\bar{P}_\lambda} \right)^2 \right] . \quad (13.48)$$

Interestingly, the noise floor is proportional to $(\bar{P}_\lambda)^{-2}$ at low power, and to $(\bar{P}_\lambda)^{-1}$ above the threshold power

$$P_{\lambda,t} = \frac{1}{2} \frac{FkT}{R} \frac{h\nu_\lambda}{q^2\eta} \quad (13.49)$$

For example, taking $\nu_\lambda = 193.4$ THz (wavelength $\lambda = 1.55 \mu\text{m}$), $\eta = 0.6$, $F = 1$ (noise-free amplifier), and $m = 1$, we get a threshold power $P_{\lambda,t} = 335 \mu\text{W}$, which sets the noise floor at $5.1 \times 10^{-15} \text{ Hz}^{-1}$ (-143 dB/Hz).

Figure 13.17 shows the scheme of a correlation system for the measurement of the microwave AM noise. It may be necessary to add a microwave amplifier at the output of each photodetector. Eq. (13.48) holds for one arm of Fig. 13.17. As there are two independent arms, the noise power is multiplied by two.

Finally, it is to be pointed out that the results of this section concern only the white noise of the photodetector and of the microwave amplifier at the photodetector output. Experimental method and some data in the close-in microwave flickering of the high-speed photodetectors is available in Reference [?]. The noise of the microwave power detector and of its amplifier is still to be added, according to Section 13.6.

13.11 Calibration

For small variations ΔP around a power P_0 , the detector gain is replaced by the differential gain

$$k_d = \frac{dv_d}{dP} \quad . \quad (13.50)$$

which can be rewritten as

$$k_d = \frac{\Delta v_d}{\frac{\Delta P}{P_0} P_0} \quad . \quad (13.51)$$

Equations (13.10)–(13.11), which are used to get $S_\alpha(f)$ from the spectrum $S_v(f)$ of the output voltage in single-channel measurements, rely upon the knowledge of the calibration factor $k_d P_0$. The separate knowledge of k_d and P_0 is not necessary because only the product $k_d P_0$ enters in Eq. (13.10)–(13.11). Therefore we can get $k_d P_0$ from

$$k_d P_0 = \frac{\Delta v_d}{\Delta P / P_0} \quad . \quad (13.52)$$

This is a fortunate outcome for the following reasons

- A variable attenuator inserted in series to the oscillator under test sets a static $\delta P/P_0$ that is the same in all the circuit; this is a consequence of linearity. For reference,

step, dB	$\Delta P/P_0$
0.1	2.33×10^{-2}
0.5	0.122
1	0.259


- A power ratio can be measured (or set) more accurately than an absolute power.

Some strategies can be followed (Fig. 13.18), depending on the available instrumentation. In all cases it is recommended to


- make sure that the power detector works in the quadratic region (see Fig. 13.5) by measuring the power at the detector input.
- exploit the differential accuracy of the instruments that measure ΔP and ΔV , instead of the absolute accuracy. Use the “relative” function if available, and do not change input range.
- avoid plugging and unplugging connectors during the measurement. A directional coupler is needed not to disconnect the power detector for the measurement of ΔP .

In Fig. 13.18 A, the internal variable attenuator of a synthesizer is used to measure $k_d P_0$. $\Delta P/P_0$ can be measured with the power meter, or obtained from the calibration of the synthesizer internal attenuator. Some modern synthesizers have a precise

A – Synthesizer



B – By-step attenuator



C – Correlation system




Figure 13.18: Calibration schemes.

attenuator that exhibit a resolution of 0.1 or 0.01 dB. In Fig. 13.18 B, a calibrated by-step attenuator is inserted between the source under test and the power detector. By-step attenuators can be accurate up to some 3–5 GHz. Beyond, one can use a multi-turn continuous attenuator and rely on the power meter. In the case of correlation measurements (Fig. 13.18 C), symmetry is exploited to measure k_a and k_b in a condition as close as possible to the final measurement of $S_\alpha(f)$. Of course, it holds that $\Delta P_a/P_a = \Delta P_b/P_b$.

13.11.1 Alternate calibration method

Another method to calibrate the power detector makes use of two synthesizers in the frequency region of interest, so that the beat note falls in the audio frequencies (Fig. 13.19). This scheme is inspired to the two-tone method, chiefly used to measure the deviation of the detector from the ideal law $v_d = k_d P$ [?, ?].

Using $P = \frac{v^2}{R}$, and denoting the carrier and the reference sideband with $v_0(t) =$

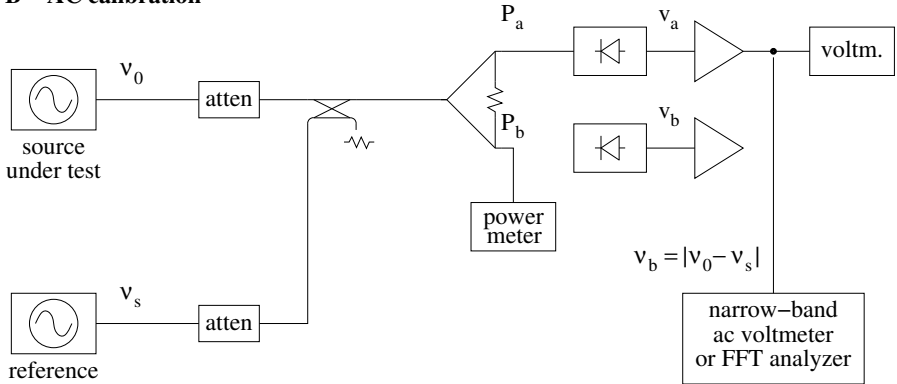
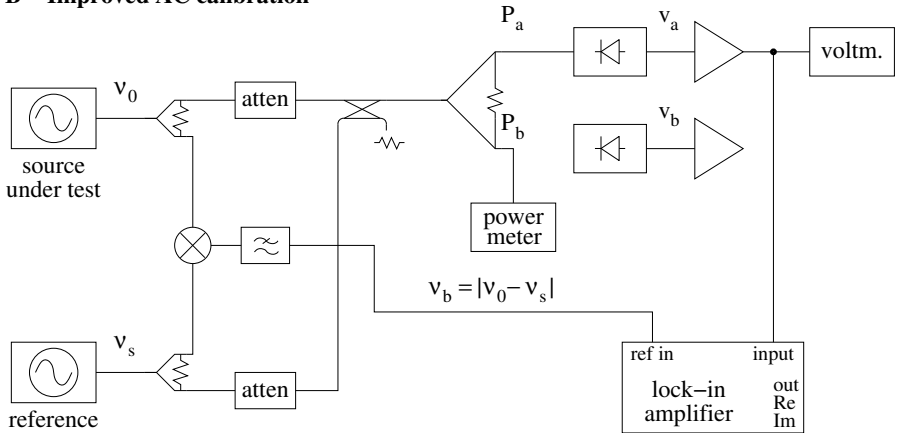
B – AC calibration**B – Improved AC calibration**

Figure 13.19: Alternate calibration schemes.

$V_0 \cos(2\pi\nu_0 t)$ and $v_s(t) = V_s \cos(2\pi\nu_s t)$, respectively, the detected signal is

$$v_d(t) = \frac{k_d}{R} \left\{ v_0(t) + v_s(t) \right\}^2 * h_{lp}(t) \quad . \quad (13.53)$$

The low-pass function h_{lp} keeps the dc and the beat note at the frequency $\nu_b = \nu_s - \nu_0$, and eliminates the $\nu_s + \nu_0$ terms. Thus,

$$v_d(t) = \frac{k_d}{R} \left\{ \frac{1}{2} V_0^2 + \frac{1}{2} V_s^2 + 2 \frac{1}{2} V_0 V_s \cos [2\pi(\nu_s - \nu_0)t] \right\} \quad , \quad (13.54)$$

which is split into the dc term

$$\bar{v}_d = k_d \frac{V_0^2 + V_s^2}{2R} \quad (13.55)$$

and the beat-note term

$$\tilde{v}_d(t) = 2 \frac{k_d}{R} \frac{V_0 V_s}{2} \cos [2\pi(\nu_s - \nu_0)t] , \quad (13.56)$$

hence

$$(V_d)_{\text{rms}} = k_d \sqrt{2P_0 P_s} . \quad (13.57)$$

The dc term [Eq. (13.55)] makes it possible to measure k_d from the contrast between \bar{v}_1 , observed with the carrier alone, and \bar{v}_2 , observed with both signals. Thus,

$$k_d = \frac{\bar{v}_2 - \bar{v}_1}{P_s} \quad (13.58)$$

Alternatively, the ac term [Eq. (13.57)] yields

$$k_d = \frac{(V_d)_{\text{rms}}}{\sqrt{2P_0 P_s}} \quad (13.59)$$

The latter is appealing because the assessment of k_d relies only on ac measurements, which are free from offset and thermal drift. On the other hand, the two-tone measurement does not provide the straight measurement of the product $k_d P_0$.

13.12 Examples


Figure 13.20 show an example⁷ of AM noise measurement. The source under test is a 100 MHz quartz oscillator (Wenzel 501-04623E serial no. 3752-0214).

Calibration is done by changing the power $P_0 = -10.2$ dBm by ± 0.1 dB. There results $k_a = 1.28 \times 10^5$ V/W and $k_b = 1.34 \times 10^5$ V/W, including the 52 dB amplifier (321 V/W and 336 V/W without amplification). The system gain is therefore $4k_a k_b P_a P_b = 641$ V² (28.1 dBV²).

The cross spectrum of Fig. 13.20 is $S_{ba} = 1.26 \times 10^{-11}$ V² (−109 dBV²/Hz) at 10 Hz, of the flicker type. Averaging over $m = 2104$ spectra, the single-channel noise is rejected by $\sqrt{2 \times 2104} = 64.9$ (18.1 dB). The displayed flicker (−109 dB at 10 Hz) exceeds by 6.4 dB the rejected single-channel noise. A correction of a factor 0.77 (−1.1 dB) is therefore necessary⁸. The corrected flicker is $S_{ba} = 9.7 \times 10^{-12}$ V² (−110.1 dBV²/Hz) extrapolated at 1 Hz. The white noise can not be obtained from Fig. 13.20 because of the insufficient number of averaged spectra.

⁷Vol. 7 p. 28, (scheme at p. 18).

⁸In the first experiments (cf. my notebooks, Vol. 7 p. 28), I did not use the correction.



Wenzel 501-04623E (s/n 3572-0214)

$P_d = 9.5 \mu\text{W}$ (-10.2 dBm)

$k_d = 1.31 \times 10^5 \text{ A}^{-1}$ with dc ampli

$m = 2104$ averaged spectra

$\sqrt{2m} = 64.9$ (18.1 dB)

$h_{-1} = 1.5 \times 10^{-13} \text{ Hz}^{-1}$ (-128.2 dB)

$\sigma_\alpha = 4.6 \times 10^{-7}$

Wenzel 501-04623E 100 MHz OXCO




Figure 13.20: Example of AM noise measurement (Vol. 7 p. 28). Gain 28.2 dB.

As a consequence of the low amplitude noise of the oscillator, it is possible to measure the noise of single channel, which includes detector and amplifier. Accounting for the gain (28.1 dBV^2), the single-channel flicker noise of Fig. 13.20 at 1 Hz is $S_\alpha(1\text{Hz}) = 2.5 \times 10^{-12} \text{ Hz}^{-1}$ (-116.1 dB/Hz) for one channel, and $S_\alpha(1\text{Hz}) = 3.4 \times 10^{-12} \text{ Hz}^{-1}$ (-114.7 dB/Hz) for the other channel.

The AM flickering of the oscillator is $S_\alpha(1\text{Hz}) = 1.15 \times 10^{-13} \text{ Hz}^{-1}$ (-129.4 dB/Hz), thus $h_{-1} = 1.15 \times 10^{-13}$. Using the conversion formula of Tab. 13.1 for flicker noise, the Allan variance is $\sigma_\alpha^2 = 1.6 \times 10^{-13}$, which indicates an amplitude

Table 13.7: AM noise of some sources.

source	h_{-1}	σ_α	notes
Anritsu MG3690A synthesizer (10 GHz)	2.5×10^{-11} −106.0 dB	5.9×10^{-6}	Fig. 13.21 A (Plot 607) (VI-140/143/147)
Marconi synthesizer (5 GHz)	1.1×10^{-12} −119.6 dB	1.2×10^{-6}	(VI-168)
Macom PLX 32-18 0.1 → 9.9 GHz multiplier	1.0×10^{-12} −120.0 dB	1.2×10^{-6}	Fig. 13.22 A (IX-133)
Omega DRV9R192-105F 9.2 GHz DRO	8.1×10^{-11} −100.9 dB	1.1×10^{-5}	Fig. 13.21 B bump and junks (Plot 630, VII-13-14)
Narda DBP-0812N733 amplifier (9.9 GHz)	2.9×10^{-11} −105.4 dB	6.3×10^{-6}	Fig. 13.22 A/B (IX-125)
HP 8662A no. 1 (CE122) synthesizer (100 MHz)	6.8×10^{-13} −121.7 dB	9.7×10^{-7}	Fig. 13.21 C junks (Plot 631, VII-17)
HP 8662A no. 2 (CE125) synthesizer (100 MHz)	1.3×10^{-12} −118.8 dB	1.4×10^{-6}	Fig. 13.21 D junks (Plot 632, VII-18)
Fluke 6160B (CE124) synthesizer	1.5×10^{-12} −118.3 dB	1.5×10^{-6}	Fig. 13.21 E junks (Plot 633, VII-19)
Racal Dana 9087B synthesizer (100 MHz)	8.4×10^{-12} −110.8 dB	3.4×10^{-6}	Fig. 13.21 F junks (Plot 634, VII-19)
Wenzel 500-02789D 100 MHz OCXO (# 0627-9308)	4.7×10^{-12} −113.3 dB	2.6×10^{-6}	Fig. 13.20 (Plot 638, VII-28)
Wenzel 501-04623E no. 1 100 MHz OCXO (# 3752-0214)	2.0×10^{-13} −127.1 dB	5.2×10^{-7}	(VII-28)
Wenzel 501-04623E no. 2 100 MHz OCXO (# 3753-0214)	1.5×10^{-13} −128.2 dB	4.6×10^{-7}	(VII-33)

stability $\sigma_\alpha = 4 \times 10^{-7}$, independent of the measurement time τ .

Table 13.7 shows some examples of AM noise measurement. The measured spectra are in Fig. 13.20, 13.21, and 13.22

All the experiments of Tab. 13.7 and Fig. 13.20, 13.21, and 13.22 were done be-







A: Anritsu MG3690A synthesizer**B: DRO Omega DRV9R192-105F****C: HP 8662A synthesizer****D: HP 8662A synthesizer****E: Fluke 6160B synthesizer****F: Racal Dana 9087B synthesizer**

Figure 13.21: Examples of AM noise measurement (Vol. 6 p. 143).

fore thinking seriously about the design of the front-end amplifier (Section 13.7), and before measuring the detector gain as a function of the load resistance (Table 13.4, and Figures 13.6–13.7). The available low-noise amplifiers, designed for other purposes, turned out to be a bad choice, far from being optimized for this application. Nonetheless, in all cases the observed cross spectrum is higher than the limit set by the average of two independent single-channel spectra. In addition, the limit set by channel isolation is significantly lower than the observed cross spectrum. These two facts indicate that the measured cross-spectrum is the true AM noise of the source. Thus Table 13.7 is an accurate database for a few specific cases. Of course, Table 13.7 also provides the order of magnitude for the AM noise of other synthesizers and oscillators employing similar technology. On the other hand, the data of Table 13.7 do




Figure 13.22: Examples of AM noise measurement.(File am-plot706salpha and am-plot705salpha.)

not provide information on the detector noise.

The amplifier used in almost all the experiments is the “super-low-noise amplifier” proposed in the data sheet of the MAT03 [Analog-MAT03, , Fig. 3], which is matched




Figure 13.23: Trans-impedance configuration for the measurement of AM noise.

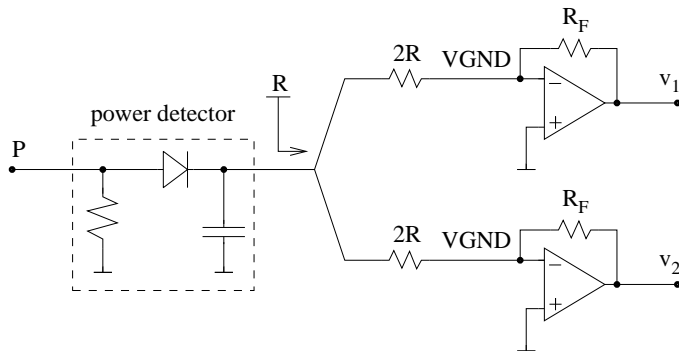


Figure 13.24: Dual-channel trans-impedance measurement of AM noise.

PNP transistor pair. For reference, the NPN version of this amplifier is discussed in [?]. The input differential stage of this amplifier consists of three differential pairs connected in parallel, so that the voltage noise is the noise of a pair divided by $\sqrt{3}$. Yet the current noise is multiplied by $\sqrt{3}$. As a consequence, the amplifier is noise-matched to an impedance of some $200\ \Omega$ for flicker noise, and to some $30\ \Omega$ for white noise, which is too low for our purposes. The second version of the MAT03 amplifier, designed after the described experiments, was optimized for the lowest flicker when connected to a $50\ \Omega$ source [?]. This amplifier, now routinely employed for the measurement of phase noise, makes use *one* MAT03 instead of three. In two cases (Fig. 13.22) a different amplifier was used, based on the OP37 operational amplifier loaded to an input resistance of some $1\ \text{k}\Omega$. Interestingly, in the operating conditions of AM noise measurements, the OP37 outperforms the more sophisticated MAT03.

13.13 Final remarks

True quadratic detection vs. peak detection. Beyond a threshold power, a power detector leaves the quadratic operations and works as a peak detector. The peak detection is the same operation mode of the old good detectors for AM broadcasting (which is actually an *envelope* modulation). This operation mode exhibits higher gain, hence it could be advantageous for the measurement of low-noise signals. The answer may

depend on the diode type, Schottky or tunnel. The strong recommendation to use the diode in the quadratic region might be wrong.

Trans-resistance amplifiers. In principle, the power detector can be used as a power-to-current converter (instead of as a power-to-voltage converter), and connected to a trans-resistance amplifier (Fig. 13.23). The advantage is that the resistor at the detector output, which is a relevant source of noise in voltage-mode measurements, is not present. On the other hand, the role of the detector internal capacitor depends on the residual input impedance of the amplifier, and ultimately on the amplifier bandwidth. At HF/VHF/low-UHF frequencies, the residual carrier present at the detector output may impair the operation of the amplifier input stage. The trans-resistance configuration, suggested in [?], is never found in the technical literature accompanying the detectors.

Dual-channel trans-resistance amplifiers. For a while I was intrigued by the scheme of Fig. 13.24, with the purpose to eliminate the resistor noise by correlation and averaging. This is not possible because the current noise of the resistor, present at the input of the two amplifiers, is fully correlated. A solution, if any, can only be based on four-port directional coupler, as suggested in [?].

Cryogenic environment. In principle, the tunnel diode should work at cryogenic temperatures. Yet, the laboratory could be much less smooth than the theory.

NIST AM noise facility. I have only little information about the NIST AM noise facility, most of which from the FCS.

In the HF-VHF band (up to 100 MHz), the noise floor of the NIST AM noise facility is of -150 dBc/Hz. Nothing is said about flicker.

Looking at the example of Sec. ??, the flicker noise of the oscillator is $S_{\alpha}(1\text{Hz}) = -129.4$ dB/Hz, i.e. -132.4 dBc/Hz at 1 Hz, thus -152.4 dBc/Hz at 100 Hz. Therefore, the example shows the measurement of an actual oscillator below -150 dBc/Hz at $f = 100$ Hz, well before the floor is reached.

Looking at Fig. 13.20, the floor of the instrument is not worse than -160 dBc/Hz. With a pinch of good luck, it could be -170 dBc/Hz with longer averaging time. Averaging on 32767 spectra instead of 638 gives a further noise reduction of 8.5 dB.

Chapter 14

Exotic Methods

ZZZZZZZZZZZZZZZ

Chapter 15

Applications

15.1 A Taste of Phase Noise and Frequency Stability

There are many reasons to be interested in phase noise and frequency instability, most of which are related with spectrum broadening, with timing uncertainty, and with reduced coherence time. The impact of phase noise and frequency instability is surprisingly ubiquitous, from everyday technology to fundamental science. Let us go through some examples.

In radio systems, sensitivity or dynamic range may be limited by phase noise sidebands appearing (i) on the local oscillator of a superheterodyne receiver, causing neighboring strong signals to mask the desired IF signal (reciprocal mixing) [Grebekemper, 1981], [Sosin, 1971, Sec. 2.5], (ii) on strong adjacent transmitted signals themselves in channelized systems (the near-far problem, in which sideband noise from local users can overpower distant cellular base stations), or (iii) on strong reflected radar signals such as ground clutter, limiting the ranging accuracy and the detection of small objects in the clutter [Skolnik, 2008, Chap. 6], [Leeson and Johnson, 1966, Leeson and Johnson, 1964]. In these systems, amplitude fluctuation (AM) is limited in oscillators and nonlinear power amplifiers, so phase noise (PM) is generally the dominant source of noise sidebands. Similar problems are found in LIDAR, RADAR's optical counterpart. Phase- and frequency-modulated systems are, of course, directly subject to noise in the frequency domain; examples include analog frequency-division FM multiplex systems and digital QAM modulation systems.

⁰Everything still to be done...

In digital systems and communications, and in microelectronics as well, the term ‘jitter’ is often used for phase noise integrated over the appropriate bandwidth and converted into time fluctuations. In turn, time fluctuations increase the bit error rate (see for example [Da Dalt and Sheikoleslami, 2018] and [Li, 2008]). In such systems, the ‘quality’ of the digital clock is often shown as the eye pattern (eye diagram), which is corrupted by the jitter (see [Da Dalt and Sheikoleslami, 2018, Sec 1.2], [Li, 2008, Sec. 4.1] and ‘Eye pattern’ entry on Wikipedia). Accurate clock synchronization is a major leap forward in 5G/6G wireless systems [Lin, 2018, Tataria et al., 2021].

The community of power grids is looking at frequency stability for future power grids [Conte et al., 2021, Kruse et al., 2021, RG-CE System Protection and Dynamics Sub Group, 2016] because sustainability requires a spread of generator technologies with different inertia (latency between demand for increased/reduced power, and actual power delivery). We believe that, at some point the “generators’ inertia” will be identified with the integration time τ of the Allan variances. Insufficient synchronization is blamed as a co-factor of the 2003 Northeast Blackout in the USA [Amelot et al., 2012].

Unique requirements arose with the development of the Deep Space Network. The DSN required a blend of long-term stability for initial acquisition of weak signals and short-term stability for maintaining lock. This new combined need led to the 1964 IEEE-NASA Symposium on Short-term Frequency Stability [Chi, 1964], to a Special Issue of Frequency Stability published in 1966 in the Proc. IEEE [Chi, 1966], and to the definitive 1971 Barnes et al. article [Barnes et al., 1971]. The latter is a precursor of the IEEE Standard 1139 [Donley et al., 2022]. That effort to update this standard continues today.

The Allan variance tools also find applications in geodesy and astrometry [Malkin, 2016]. The short-term fluctuations of oscillators used in Very-Large Baseline Interferometry (VLBI) cannot be compensated numerically, and limit the detection sensitivity. Time fluctuations are critical in gravitational wave experiments (LIGO) because of the tiny spacetime warp to be detected, and in the RF cavities of particle accelerators [Holzer, 2018, Chapters by F. Tecker, and H. Damereau], where they cause intensity loss or energy loss in the beam.

In quantum computing, the correlation time limits the lifetime of a qubit [Ball et al., 2016].

Optics supports a major trend in high-purity signals [Diddams et al., 2020], and in atomic oscillators and clocks as well [Udem et al., 2002], because the time interval associated with a phase angle is $\sim 10^{-4}$ smaller than at microwave frequencies. The femtosecond laser, which enabled the first direct synthesis from RF to optics, played a major rôle in this trend. The Nobel Prize in Physics was awarded to John L. Hall and Theodor W. Hänsch in 2005 for the femtosecond laser (together with Roy J. Glauber). See also [Cundiff and Ye, 2003, Kuse and Fermann, 2017, Xie et al., 2017]. Phase noise challenges the networks intended for fundamental science and clock comparison between metrological labs using optical fibers shared Internet data traffic [Calosso et al., 2016, Lisdar et al., 2016, Delva et al., 2017, Nicolodi et al., 2014].

In the new version of the International System of Units SI [Bureau International des Poids et Mesures (BIPM), 2019], virtually all measurement units rely on time. Thus time fluctuations, and equivalently frequency fluctuations, impact all branches

of metrology. For example, the ampere, unit of electric current, relies on counting the electrons flowing in the unit of time. And the Josephson voltage standard is based on the conversion from photon energy to voltage. The Nobel Prize in Physics was awarded to Brian D. Josephson in 1973 for the theoretical prediction of this effect.

15.2 Radars and radioastronomy¹

Yesterday Paul Boven explicated the impact of phase noise of the local oscillator on the coherent accumulation of energy in a correlation. He provides the reference <https://link.springer.com/book/10.1007/978-3-319-44431-4> page 434 – where an analytical solution on the degradation of SNR as a function of phase noise ... fascinating. Might provide an answer finally as to why high quality RADAR are looking for good LO.

¹Jean-Michel email, March 31, 2023

Chapter 16

The companion of the Enrico's Chart

16.1 Introduction

Phase noise and frequency instability are equivalent concepts, to the extent that frequency is the derivative of the instantaneous phase. However, the choice of terminology, of mathematical tools and of experimental methods really depends on what aspect of the noise our system is sensitive to, or which aspect we focus on. The Power Spectral Density (PSD) of the phase fluctuations of a periodic signal, as a function of the Fourier frequency, is the preferred tool to describe phase noise. Similarly, the frequency instability is described by the PSD of frequency fluctuations, or more often by the two-sample (Allan or Allan-like) variance of the fractional frequency, as a function of the measurement time.



Download
Enrico's Chart
from Zenodo

Download this
article from
IEEE Xplore



Figure 16.1: QR codes to download Enrico's Chart and this article.

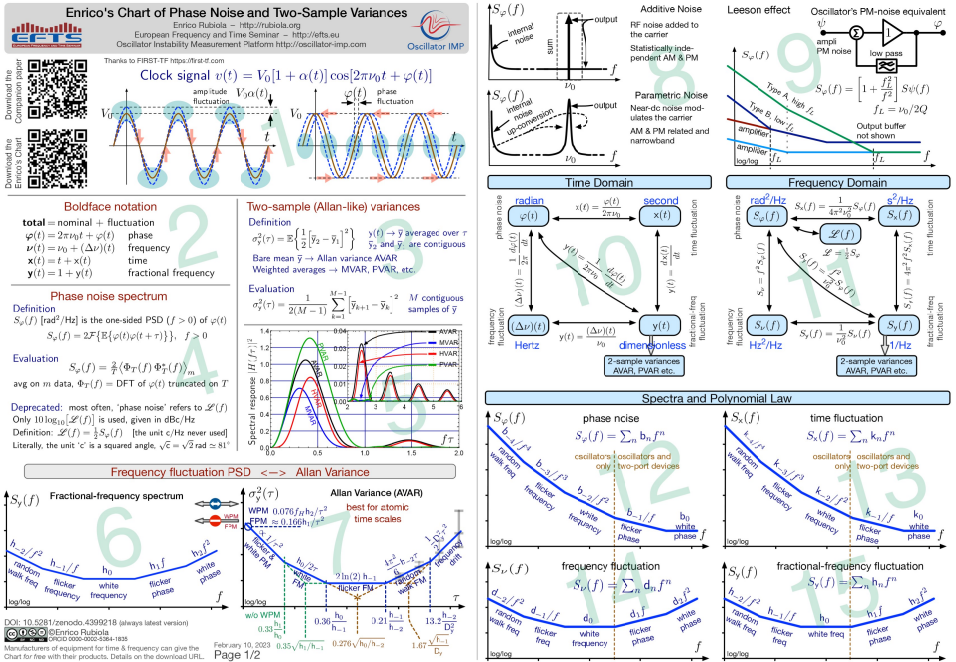


Figure 16.2: Enrico's Chart of Phase Noise and Two-Sample Variances, front side.

A Taste of Phase Noise and Frequency Stability

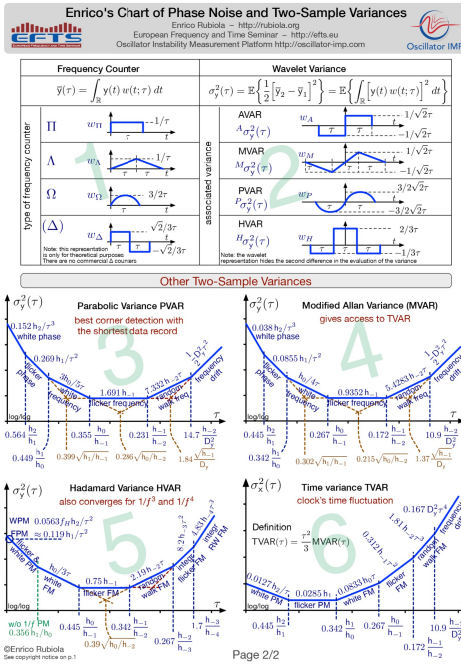
16.1.1 Enrico's Chart

Enrico's Chart of Phase Noise and Two-Sample Variances is a reference card collecting the most useful definitions, formulas and plots in the domain of phase noise and frequency stability. Starting from the first draft in 2011, it has been improved from a loose single page on the back of the program of the European Frequency and Time Seminar¹ to quite a dense front/back format. At least 1000 plasticized prints have been distributed as learning material at the EFTS and at invited conferences, courses and seminars. The current form results from the scientific contribution of the two authors, and from feedback and amendments by the users. The name on the Chart is mainly historical, but it reflects Enrico's tenacity in maintaining, updating, improving the graphic design, and in distributing the Chart.

16.1.2 License and Distribution

Enrico's Chart is a digital object (PDF file) available from Zenodo as DOI 10.5281/zenodo.4399218, and released under Creative Commons 4.0 CC-BY-NC-ND license. This DOI is called *concept DOI* because it always resolves to the latest ver-

¹The EFTS <http://efts.eu> is a full-week crash course with lectures and lab sessions funded in 2013, and held in Besançon every summer since.



noise type	$S_y(f)$	Spectra to Variances Conversion					noise type
		AVAR $\sigma_y^2(\tau)$	MVAR $\sigma_y^2(\tau)$	HVAR $\sigma_y^2(\tau)$	PVAR $\sigma_y^2(\tau)$	TVAR $\sigma_y^2(\tau)$	
Blue FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
White FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
Pink FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
White PM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
Pink PM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
Random walk FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
Integrated FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
Integrated BW FM	$\frac{3}{2}f_0^2$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	$\frac{3}{2}f_0^2 \ln(8 + 10 \ln(f_0/f)) \frac{h_1}{h_2}$	
linear drift D_0							

Figure 16.3: Enrico's Chart of Phase Noise and Two-Sample Variances, back side.

sion. Albeit Zenodo delivers a separate DOI for each version, version DOIs should only be used in special cases. The QR codes of Fig. 16.1 point to the Chart and to this article. Redistribution is encouraged as the Internet link or as the QR code, not as a file.

The 'NC' and 'ND' copyright attributions deserve a comment. The 'NC' restriction is mainly intended to prevent selling the Chart for profit. In contrast, if Companies distribute the Chart for free with their products, that is *not considered commercial use*, and we *encourage* this practice. The equipment employed for measurements in time and frequency is a notable example. The 'ND' restriction preserves the academic independence of the Chart, preventing Companies from modifying it, or from building advertisements upon it.

For any other use, Enrico himself is the preferred contact.

16.1.3 About this Article

This article is intended (i) to accompany the Chart, (ii) as a short tutorial, and (iii) as learning material for lectures. Besides, (iv) it can be cited as a summary of modern notation, provided the reader agrees with our choices.

The reader having an already good understanding of time and frequency may go straight to Enrico's Chart and put this article aside for later reading, or refer to

students and younger colleagues. Less experienced people may appreciate this article as a tutorial or as a review. To them, if they have been around for long enough, phase noise and frequency stability may relate to a well identified problem.

A reduced copy of the Chart is included (Fig. 16.2 and 16.3); the indication “*Region 1.n*” in the text means that we refer to the region n defined by the watermarks on page 1 (Fig. 16.2) of the Chart, while “*Region 2.n*” refers to the region n on page 2 (Fig. 16.2). In any event, it is a good idea to have on hand a separate copy of the Chart.

16.1.4 Suggested Introductory Readings

We advise starting from the latest version of the IEEE Standard 1139 [Donley et al., 2022]. The appendices of this standard cover topics similar to ours, but presentation and standpoint are surprisingly different. This article is our personal view, whereas the IEEE Standard is the outcome of a rather large committee. The Recommendations of the International Telecommunication Union (ITU) [ITU, 2017, ITU Working Group 15, 2000] are a must too. The European Telecommunications Standards Institute (ETSI) provides a series of standards related to the clock synchronization in networks [ETSI, 2002, Part 1-7].

Barnes et al. 1971 [Barnes et al., 1971] is one of the very first articles that all readers should study. It defines the language and the notation still in use, it introduces the spectra, the variances and their relations, and it explains the early experimental methods. The quantities φ , x , $\Delta\nu$ and y were first defined there. People interested in the birth of the scientific ideas that originated most concepts of phase noise and of frequency stability should review [Chi, 1964]. This is the Proceedings book of a one-time IEEE-NASA conference that took impetus from the DSN. Rutman 1978 [Rutman, 1978] is a review article about the progress on the concepts of phase noise and frequency stability after the early ideas. See also [Rutman and Walls, 1991].

Turning our attention to books and booklets, the Riley handbook [Riley,] is free available, sponsored by NIST. Emphasis is on Allan and Allan-like variances, rather than noise spectra, making extensive use of the Stable32 software package². Owen 2004 [Owen, 2004] is a rather extensive practical guide about phase noise, albeit elderly. Kroupa 1983 [Kroupa, 1983] is a collection of classical articles about phase noise and frequency stability, a few of which are cited elsewhere in this article. The same author also published a monograph [Kroupa, 2012]. Reference [Sullivan et al., 1990] is another edited book collecting classical articles from NIST. Rohde et al. 2021, [Rohde et al., 2021, Chapter 2 (136 pages)] is a recent text entirely about phase noise and frequency stability.

Additional material and large slideshows are available on the Enrico's home page <http://rubiola.org>. Upon request, the original PPTX files may be released to qualified users. We turn now to the Chart and Fig. 16.2.

²This and other software packages are discussed in Section 4.3.

Phase Noise and Amplitude Noise

16.1.5 The Clock Signal (Region 1.1)

Phase Noise Spectrum (Region 1.4)

Deeper Thoughts About the PSD

The Quantity $\mathcal{L}(f)$ and the Related Measurement Units

The Former Definition of $\mathcal{L}(f)$, and the Deprecated Terms ‘SSB Noise’ and ‘Offset Frequency’

Suggested Readings

Useful Quantities (Region 1.10–1.11)

Phase Time Fluctuation (or Phase Time) $x(t)$

Frequency Fluctuation $(\Delta\nu)(t)$

Fractional Frequency Fluctuation $y(t)$

The polynomial law, or power law

The Quantities φ , x , $\Delta\nu$ and y in Frequency Synthesis

Digital dividers

Output stage

High-order multiplication

Phase-Type (φ -Type) and Time-Type (x -Type) PM Noise

Phase-Type (φ -Type) Phase Noise

Time-Type (x -Type) Phase Noise

Aliasing

Notation

16.2 Two-Port Components (Region 1.8)

Additive and Parametric Noise

Added Noise

White Noise

Flicker Noise


Input-to-Output Delay

16.3 Oscillators and the Leeson Effect (Region 1.9)

The oscillator (Fig. 16.4a) consists of a loop where the resonator sets the oscillation frequency ν_0 and the sustaining amplifier compensates for the resonator loss. Gain clipping (nonlinearity) is necessary to stabilize the amplitude, hence the typical dominance of PM over AM. The buffer isolates the loop from the load.

Everyday experience suggests that, unlike the two-port components, the time fluctuation $x(t)$ of an oscillator can be quite large. The reason is that the oscillator

(a)



(b)




Figure 16.4: Phase noise mechanisms inside an oscillator. (a) Conceptual schematic of the oscillator, where signal is the microwave/RF sinusoidal oscillation at the frequency ν_0 . (b) Phase-noise equivalent of the oscillator, where the signal is the random phase of the microwave/RF circuit shown above. In this representation, phase noise is always additive, regardless of its physical origin. The grayed blocks in (a) and (b) represent the noise of the resonator and of the buffer, not included in equations.

accumulates the ‘error’ of each cycle, however small it may be.

Resonator, sustaining amplifier and buffer, they all introduce phase noise (amplitude noise too, but this is a more specialized topic, not considered here). The phase modulators in Fig. 16.4a are necessary to model phase noise of the various part of the oscillator in the general case, including white and colored processes. However, these modulators make the mathematical treatise difficult. The *phase equivalent* model, introduced in [Rubiola, 2010, Chapter 4] and shown in Fig. 16.4b, avoids such difficulty. The signal circulating in Fig. 16.4b is the phase fluctuation of Fig. 16.4a. In this representation:

- Sustaining amplifier and buffer map into ‘phase amplifiers’ of gain exactly equal one because they ‘copy’ the input phase to the output, unchanged. Otherwise stated, the amplifier’s delay cannot be stretched/shrunk,
- Random phase modulation maps into noise *added* in the signal path, regardless of the nature of the process,
- The phase model is linear because gain clipping has no effect on phase—at least, not first-order effects,

- The resonator (narrow-band 2nd-order filter) maps into a single-pole low-pass filter

This representation holds for large quality factor, say, $Q \gtrsim 50$, which ensures that signals are sinusoidal. Accepting this limitation, the approximation in Fig. 16.4b adequately models the more general oscillator in Fig. 16.4a. Breaking the high- Q hypothesis, the simple mathematical treatise provided below is no longer valid because distortion introduces coupling between amplitude and phase.

We give a short summary of the original derivation [Rubiola, 2010, Sec. 4.4–4.5]. The low-pass impulse response³ of the resonator is $b(t) = (1/\tau)e^{-t/\tau}$, where $\tau = Q/\pi\nu_n$ is the relaxation time, Q is the quality factor in actual load conditions, and ν_n is the natural frequency. In practice, the oscillation frequency ν_0 is so close to ν_n that they are interchangeable. The Laplace transform of $b(t)$ is $B(s) = (1/\tau)/(s + 1/\tau)$. The cutoff frequency of the low-pass equivalent is called *Leeson frequency*, equal to half the resonator's RF bandwidth ν_0/Q

$$f_L = \frac{\nu_0}{2Q}. \quad (16.1)$$

Letting aside the buffer noise, we introduce the phase-noise transfer function of the oscillator

$$\text{definition:} \quad H(s) = \frac{\Phi(s)}{\Psi(s)}, \quad (16.2)$$

where $\Phi(s)$ and $\Psi(s)$ are the Laplace transforms of the output $\varphi(t)$ and of the input $\psi(t)$, respectively. Inspecting the loop on the left of Fig. 16.4b, we find $H(s) = 1/[1 - B(s)]$. The latter expression is derived with the classical rules of feedback, the same used in textbooks to analyze a simple control loop, of to derive the gain of the operational amplifier in non-inverting configuration. Expanding $H(s)$ and replacing $s \rightarrow j2\pi f$, we get $|H(f)|^2 = 1 + f_L^2/f^2$. The multiplication by $1/f^2$ which occurs at $f < f_L$ is the 'soul' of the Leeson effect: the oscillator integrates the phase fluctuation $\psi(t)$ occurring in the loop. The oscillator phase noise is $S_\varphi(f) = |H(f)|^2 S_\psi(f)$, thus


$$S_\varphi(f) = \left[1 + \frac{f_L^2}{f^2}\right] S_\psi(f). \quad (16.3)$$

Plugging the amplifier noise $S_\psi(f) = \mathbf{b}_{-1}/f + \mathbf{b}_0$ in (16.3), we find the two typical patterns of Fig. 16.5. Interestingly, the oscillator spectrum contains either $1/f^2$ or $1/f$ noise, not both. The $1/f^2$ noise is common in microwave oscillators, characterized by high ν_0 and low Q , thus $f_L \gg f_c$. The $1/f$ noise is typical of RF and quartz oscillators, characterized by low ν_0 and high Q .

The noise of the resonator and of the buffer, not included in (16.3) and not shown in Fig. 16.5, adds. In most quartz oscillators, the $1/f^3$ noise due to the $1/f$ fluctuation of the resonator natural frequency exceeds the $1/f$ contribution of the electronics turned into $1/f^3$ via the Leeson effect [Rubiola and Giordano, 2007]. In such cases, the Leeson effect is hidden below the resonator's fluctuations.

³The standard symbol of the impulse response is $h(t)$, and $H(s)$ its Laplace transform. Here we use $b(t) \leftrightarrow B(s)$ to save $h(t) \leftrightarrow H(f)$ for the oscillator.

(a)



(b)




Figure 16.5: Typical noise spectra occurring in oscillators. (a) Type A phase noise spectrum, often found in microwave oscillators. (b) Type B phase noise spectrum, often found in RF and quartz oscillators. Both (a) and (b) show the noise $S_\psi(f)$ of the sustaining amplifier, the flicker corner at frequency f_c , the Leeson corner at frequency f_L , and the oscillator noise $S_\varphi(f)$. The fluctuation of the resonator's natural frequency and the noise of the output buffer are not included.

16.3.1 Example of Noise in a Microwave Oscillator

Fig. 16.6 shows the phase noise of the DRO-100, a commercial 10 GHz oscillator based on a dielectric resonator. By comparison with Fig. 16.5, this oscillator is clearly of type A, characterized by $f_c < f_L$. The coefficients b_0 , b_{-2} and b_{-3} are obtained from graphical approximation of the measured spectrum with the theoretical straight line having slope 0, -2 and -3 . The frequencies f_c and f_L are obtained graphically from the straight-line approximation. We start our interpretation from the right-hand side ($f = 10$ MHz) to the left.

The white PM noise $b_0 = 10^{-17}$ rad²/Hz suggests that the power at the input of the sustaining amplifier (the location where the power is the lowest) is $P \approx 0.5$ mW. The power in the resonator is substantially the same. This result is obtained using




Figure 16.6: Example of phase noise of an oscillator. The original plot is courtesy of Ulrich L. Rohde, ©Synergy Microwave Corp., used with permission. We added the polynomial approximation (the units of b_i coefficients are omitted), the estimated amplifier noise, and other comments.

(1.30), guessing that the noise figure of such amplifier is of 1 dB, quite a plausible value for good amplifiers.

The white frequency noise b_{-2}/f crosses the white phase noise b_0 at 3.75 MHz. This is clearly the Leeson frequency f_L . Using (16.1), we estimate the resonator's quality factor $Q = 1330$ in actual load conditions.

Finally, the flicker FM noise $b_{-3} = 14.1 \text{ rad}^2/\text{Hz}^2$ crosses the white FM noise at $f_c = 100 \text{ kHz}$. Equation (16.3) with $f \ll f_c$ suggests that the flicker PM of the sustaining amplifier is of 10^{-12} rad^2 (-120 dB rad^2).

16.3.2 Suggested Readings About Oscillators

Our presentation is based on [Rubiola, 2010], which includes the theoretical proof of the Leeson effect, the analysis of noise sources in the loop and at the output, and a chapter about reverse engineering of oscillators from phase noise. Rather than focusing on the schematic, [Rubiola, 2010] analyzes the oscillator as a system. A white paper [Rubiola and Brendel, 2010] extends the theory to amplitude noise. The limitation of this approach is that it is suitable only to high- Q resonator (say, $Q \gtrsim 50$), where all relevant signals are sinusoidal.

Numerous RF and quartz oscillator schemes make use of a negative-resistance amplifier in parallel to the resonator. We propose two classical references about the noise of such amplifiers, Penfield 1960 [Penfield, 1960] and van der Ziel 1962 [van der Ziel, 1962].

Other approaches deserve attention, chiefly: (i) Models based on the limit cycle, derived from the Einstein's diffusion theory, i.e., Loh & al. 2013 [Loh et al., 2013a, Loh et al., 2013b], and Demir & al. 2000 [Demir et al., 2000], or from the Langevin equations, Kärtner 1990 [Kärtner, 1990]; and (ii) the Impulse Sensitivity Function proposed by Lee & Hajimiri [Hajimiri and Lee, 1998, Lee and Hajimiri, 2000], which is good at describing low- Q oscillators, like those commonly found in microelectronics. Additionally, Pankratz 2014 [Pankratz and Sánchez-Sinencio, 2014] provides a large survey specific to oscillators in integrated circuits.

Digging in the origins, the widely-cited van der Pol oscillator [van der Pol and van der Mark, 1927] is about chaos in oscillations, rather than phase noise. Edson 1960 [Edson, 1960] is arguably the first article that analyzes the phase noise in electronic oscillators, and Leeson 1966 [Leeson, 1966] is the article that introduces the phase noise mechanism in feedback oscillators, later known as the Leeson effect. A review article by D. B. Leeson is available [Leeson, 2016].

16.4 The Allan Variance

Definition and Evaluation

Spectral Response

Overlapped Allan Variance

16.5 Other Options for the Two-Sample Variance

16.5.1 Frequency Counters and Weighted Averages (Region 2.1)

16.5.2 Generalized Two-Sample Variances (Region 2.2)

16.5.3 The Time Variance TVAR (Region 2.6)

16.5.4 Additional Options

The Total Variance

Th_o1, Th_oH and Th_oBR

16.5.5 Choosing the Most Appropriate Variance

Normalization

AVAR

MVAR

HVAR

PVAR

16.5.6 Some Pieces of Advice

16.5.7 Example of MDEV

16.5.8 Suggested Readings About Variances

General references

Π and Λ counters, and the related statistics

The Ω counter and the Parabolic Variance

Aliasing

16.6 Relations Between Phase Noise and Variance

16.6.1 Visual Inspection of Plots (Regions 1.6–1.7)

16.6.2 Conversion from PM Noise to Allan Variance (Region 2.7)

16.6.3 Variance to Spectrum Conversion (Regions 1.6, 1.7, 2.7)

16.6.4 The Cutoff Frequency f_H and the Sampling Interval τ_0 16.6.5 Example (f_H)

First option

Second option

16.7 Confidence Intervals

16.7.1 Direct Problem

16.7.2 Inverse Problem

16.7.3 Application of the Bayesian Statistics to Allan Variances

16.8 Measurement Techniques

16.8.1 Mathematical Tools for Spectral Analysis

The estimation of power spectra based on the Fast Fourier Transform (FFT) is common to all phase noise analyzers, analog or digital. The FFT is a fast algorithm for the evaluation of the Discrete Fourier Transform (DFT), and gives the same result. The algorithm used in most software packages and libraries is the ‘Fastest Fourier Transform of the West’ FFTW3 [Frigo and Johnson, 2005]. With a bit of humor, we note that the FFTW comes from the Eastern USA. The book [Brigham, 1988] is Enrico’s favorite reference about the FFT, while [Babitch and Oliverio, 1974] is about the early use of the FFT in phase noise measurements.

The most commonly used algorithm for the estimation of the PSD, introduced in [Welch, 1967] and known as the Welch algorithm, stands on the FFT. In several software packages and libraries, the PSD function takes a name which recalls Welch. This is the case of Matlab, Octave and Python (SciPy). The classical article [Harris, 1978] discusses extensively the windows (taper) functions used in spectral analysis.

Reference [Bennett, 1948] introduces the spectral analysis of quantized signal, and [Widrow et al., 1996, Widrow and Kollár, 2008] provide an extensive treatise of quantization noise.

16.8.2 How the Background-Noise Rejection Works

The noise rejection is based on the *cross-spectrum method*, described below. Referring to Fig. 9.1, we define

$$x = \varphi - \theta \quad \leftrightarrow \quad X = \Phi - \Theta \quad (16.4)$$

$$y = \varphi - \psi \quad \leftrightarrow \quad Y = \Phi - \Psi, \quad (16.5)$$

where ‘ \leftrightarrow ’ stands for Fourier transform inverse-transform pair, time and frequency are implied, and x and y should not be mistaken for x and y . For us, φ is the *signal*, θ and ψ are the *noise* to be rejected. The quantities θ and ψ account for all the noise sources, i.e., references, synthesizers, mixers, low-noise amplifiers, etc.

The cross spectrum averaged on m acquisitions is

$$S_{yx} = \frac{2}{T} \left\langle Y X^* \right\rangle_m \quad (16.6)$$

$$= \frac{2}{T} \left\langle (\Phi - \Psi) (\Phi - \Theta)^* \right\rangle_m \quad (16.7)$$

$$= \frac{2}{T} \left\langle \Phi \Phi^* - \Phi \Theta^* - \Psi \Phi^* + \Psi \Theta^* \right\rangle_m. \quad (16.8)$$

It is seen on Fig. 9.1 that φ , θ and ψ are statistically independent (separate and independent hardware), and likewise Φ , Θ and Ψ . Thus, the mathematical expectation of $\Phi \Theta^*$, $\Psi \Phi^*$ and $\Psi \Theta^*$ is zero, and

$$S_{yx} \rightarrow \frac{2}{T} \left\langle \Phi \Phi^* \right\rangle_m \equiv S_\varphi \quad \text{for large } m, \quad (16.9)$$

which is the same as (1.17). Averaging on m acquisition, the background noise of the instrument is rejected by a factor of $\sim 1/\sqrt{m}$, i.e., 5 dB per factor-of-ten in m . The actual rejection depends on the estimator.

Most phase noise analyzers use the estimator

$$\widehat{S_{yx}} = \frac{2}{T} \left| \left\langle Y X^* \right\rangle_m \right|. \quad (16.10)$$

This estimator is biased and suboptimal [Rubiola and Vernotte, 2010]. The residual noise terms $\langle \Phi \Theta^* \rangle_m$, $\langle \Psi \Phi^* \rangle_m$ and $\langle \Psi \Theta^* \rangle_m$ are proportional to $1/\sqrt{m}$.

Reference [Rubiola and Vernotte, 2010] shows that the real-part estimator

$$\widehat{S_{yx}} = \frac{2}{T} \Re \left\{ \left\langle Y X^* \right\rangle_m \right\} \quad (16.11)$$

is the best option. First, it follows from (16.8) that $\Phi \Phi^* \in \mathbb{R}$, the other terms are complex, and the background noise is equally split between $\Re\{Y X^*\}$ and $\Im\{Y X^*\}$. Thus, (16.11) is obviously advantageous vs. (16.10) because it keeps the entire signal and discards the unnecessary noise in $\Im\{Y X^*\}$. The residual single-channel is proportional to $1/\sqrt{2m}$, improving by a factor of 2. Otherwise stated, the measurement time for a given noise rejection is shorter by a factor of 4. A further advantage of (16.11) is that it is unbiased.

Sadly, (16.10) is often the one and only option, and when (16.11) is available, it is not the default.

The cross-spectrum method derives from radioastronomy [Hanbury Brown et al., 1952] and from the early attempts to assess the frequency fluctuations of Hydrogen masers [Vessot et al., 1964]. Reference [Walls et al., 1976, Section IV] is arguably the first application to the measurement of phase noise. It goes without saying that the

cross-spectrum method is of broader usefulness than just phase noise. See [Rubiola and Vernotte, 2010] for a tutorial, [Baudiquez et al., 2020] for the use of Bayesian statistics, needed to assess the statistical uncertainty when m is small, and [Baudiquez et al., 2022] for the extension to multiple instruments simultaneously measuring the same physical quantity.

16.8.3 Limitations of the Cross-Spectrum Method

The noise limit $\sim 1/\sqrt{m}$ is often shown on the analyzers' screen under fancy names like 'correlation factor,' 'number of correlations,' or 'Xcorr,' together with $\mathcal{L}(f)$. However, such noise rejection is useful only if the correlated disturbances are sufficiently small. In fact, a disturbing term $\delta \leftrightarrow \Delta$ introduced in (16.4)-(16.5) changes the instrument readout from S_φ to $S_\varphi + \varsigma S_\delta$, where $\varsigma = \pm 1$ is a sign which depends on whether δ has same sign or opposite sign in the two channels. This is a systematic error of unknown sign ς . Consequently, the common belief that the background noise of the instrument always results in the *over-estimation* of the DUT noise is *untrue*.

The obvious benefit of large m is high sensitivity⁴, i.e., high capability to measure small S_φ . The inevitable drawback is that a proportionally small S_δ can spoil the measurement. A gross error occurs when $S_\delta \approx S_\varphi$, but $S_\varphi + \varsigma S_\delta > 0$ still holds. Worse, if $S_\varphi + \varsigma S_\delta < 0$, the result is a blatant nonsense.

Among the causes for the disturbances δ , we mention the thermal energy in the input power splitter, the effect of AM on the mixers, and the RF crosstalk inside the instrument.

Most power splitters are either 4-port directional couplers terminated at one port, or Y resistive network. Such splitters, inherently, add an amount of anticorrelated noise originated from the thermal energy associated to the internal dissipation. Thus, $S_\varphi < kT/P$ (thermal energy divided by the carrier power) at the splitter output is a physically legitimate outcome, but it goes with a measurement error. The experimental evidence and the full theoretical proof is found in [Rubiola and Giordano, 2000]. The splitter's thermal noise may result in the collapse of the cross spectrum in the measurement of oscillators [Nelson et al., 2014]. Different options for the power splitter are discussed in [?]. Another perspective on the power splitters is proposed in [Gruson et al., 2017], with original results. In [Gruson et al., 2020] we propose a method for the measurement of the bias error due to both power splitter and internal crosstalk.

The correction for the splitter's thermal energy is trivial if the instrument measures the carrier power. NoiseXT/Arcale informed us that this correction now has been implemented in the DNA⁵.

16.8.4 Metrologist's Perspective of Phase Noise Uncertainty

The essential concepts we need to understand uncertainty are defined by the International Vocabulary of Metrology [Joint Committee for Guides in Metrology (JCGM),

⁴Strictly speaking, this is an improper use of the term sensitivity.

⁵Informal discussion with L. Adrien at the IEEE EFTF IFCS Joint Meeting, Paris, April 2022.

2012], generally known as VIM. The VIM is published by the BIPM, the international organization in charge to ensure worldwide unification of measurements⁶.

The statistical noise limit $\sim 1/\sqrt{m}$ falls into *Type A evaluation of measurement uncertainty* [Joint Committee for Guides in Metrology (JCGM), 2012, entry 2.28].

Other components of uncertainty (crosstalk, AM noise, thermal energy, etc.) cannot be calculated from the measurement outcomes. They fall into *Type B evaluation of measurement uncertainty* [Joint Committee for Guides in Metrology (JCGM), 2012, entry 2.29], which relies on the analysis of the system.

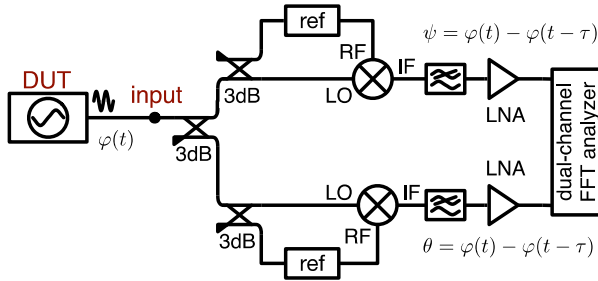
Literature of the past 10 years [Nelson et al., 2014, ?, Gruson et al., 2017, Gruson et al., 2020] and three workshops [L(f)-2014, 2014, L(f)-2015, 2015, L(f)-2017, 2017] point to substantial errors and discrepancies in the measurement of commercial oscillators exhibiting very-low phase noise. The concept of *null measurement uncertainty* applies [Joint Committee for Guides in Metrology (JCGM), 2012, entry 4.29], which is the uncertainty in the special case of signal approaching zero. Broadly speaking, this is the extrapolation of uncertainty to the measurement of a noise-free oscillator, and tells us the minimum detectable amount of phase noise.

We encourage the reader to get awareness of the *definitional uncertainty* [Joint Committee for Guides in Metrology (JCGM), 2012, entry 2.27], the component of uncertainty resulting from the lack of knowledge of details in the definition of the DUT; and of the *influence quantity*, a quantity affecting the relation between the indication and the measurement result. However, the direct application of these concepts to phase noise is not straightforward.

Reading the instruments' specs, we often find the *uncertainty* (or improperly, *error*), and the *sensitivity* parameters. For example, we may read that the uncertainty is ± 3 dB for $f < 1$ kHz, and ± 2 dB for $f \geq 1$ kHz. Similar information may be more detailed, in the form of a table. Forgiving the ' \pm ' symbol, we interpret this as the type B uncertainty. The sensitivity takes the form of a table or a spectrum, indicating $\mathcal{L}(f)$ for relevant values of ν_0 and f ; for example, -160 dBc/Hz at 1 GHz carrier and 100 kHz 'offset' (actually, modulation frequency). We are inclined to interpret this parameter as the 'null measurement uncertainty.' Note that the term 'sensitivity' found in this context is quite different from the VIM definition [Joint Committee for Guides in Metrology (JCGM), 2012, entry 4.12]. That said, we observed that users tend to trust a measured noise level below the 'sensitivity' of the instrument, provided it is still higher than the ' $\sim 1/\sqrt{m}$ ' statistical limit seen on the screen. This practice is encouraged by ads showing a 'typical' background noise, clearly lower than specs.

⁶The international coordination of metrology is a complex topic even for specialists because it has both scientific and political implications under the *Metre Convention*. The reader interested can find all information on the BIPM site <https://bipm.org>.

(a)



(b)




Figure 16.7: The discriminator method for the measurement of phase noise in oscillators. (a) Block diagram of the two-channel instrument. (b) Phase space representation of one channel.

16.8.5 The Beat Method

The method, shown in Fig. 4.4, makes use of the leverage effect which results from beating ν_0 down to $\nu_b = \nu_0 - \nu_{\text{ref}}$. Mathematically,

$$2 \cos[2\pi\nu_0 t + \varphi(t)] \cos[2\pi\nu_{\text{ref}} t + \psi(t)] = \cos[2\pi\nu_b t + \varphi(t) - \psi(t)] \quad (16.12)$$

after deleting the $\nu_b + \nu_{\text{ref}}$ term with the lowpass filter. The beat note preserves the frequency fluctuations $\nu_0 - \nu_{\text{ref}}$ and the phase fluctuation $\varphi - \psi$, and stretches the time fluctuation by a factor $\kappa \simeq \nu_0/\nu_b$.

This method is often used to beat microwave signals down to the RF region, where we can use digital instruments. In optics, this is the preferred option to measure the fluctuations of a stabilized laser with electrical instruments.

16.8.6 The Discriminator Method

The use of a reference discriminator, either a resonator or a delay line, is another way to measure the phase noise of an oscillator by comparing the oscillator output to a delayed version of the same signal.

Fig. 16.7 shows the principle and the equivalent scheme in phase space. The latter follows the same approach used with the oscillator (Sec. 16.3). Using the upper case for the Fourier transform of the lower-case function of time, the output is $\Psi(f) =$

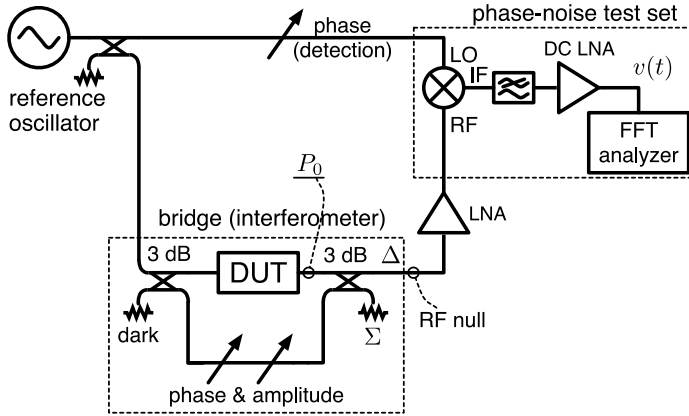


Figure 16.8: Scheme of the bridge (interferometric) phase noise measurement.

$k_d [1 - H(f)] \Phi(f)$, where the discriminator's phase response is $H(f) = (1/\tau)/(jf + 1/\tau)$ for a resonator characterized by the relaxation time τ , and $H(f) = e^{-j2\pi\tau f}$ for a delay line characterized by the delay τ . Accordingly, $\Phi(f)$ is evaluated from the analyzer readout as

$$\Phi(f) = \frac{1}{k_d [1 - H(f)]} \Psi(f) \quad (16.13)$$

Finally, using the two-channel configuration, the oscillator phase noise is evaluated as

$$S_\varphi(f) = \frac{1}{k_d^2 |1 - H(f)|^2} S_{\psi\theta}(f). \quad (16.14)$$

This idea is for sure old, albeit we could not track the origin. The scheme (Fig. 16.7) is suitable to implementations based on most or all the phase noise analyzers described in this Section. The standard 1.55- μm optical fiber of appropriate length carrying a modulated beam is a good idea to delay a microwave signal. We implemented a single-channel version [Rubiola et al., 2005], and we further developed this concept adding the cross spectrum method [Salik et al., 2004, Volyanskiy et al., 2008]. To our knowledge, the OEwaves OE8000 is the one and only commercial phase noise analyzer based on the photonic delay.

16.8.7 The Bridge (Interferometric) Method

The bridge (Fig. 16.8) is an advanced method exhibiting extremely low background noise. The bridge must be at the equilibrium, achieved by setting the bridge's phase and amplitude. At the equilibrium, the carrier is suppressed at the Δ port, and all the DUT power P_0 goes to the Σ port. The DUT noise sidebands, equally split between Σ and Δ , are amplified and synchronously detected by the mixer. The variable phase

at the mixer LO port sets the detection of PM, AM, or a combination of. Unlike the previous schemes, the mixer works in linear regime, where the RF input is a small signal, and of course the LO input is saturated.

Low background flicker b_{-1} is the main benefit. In fact, (i) the passive components used in the bridge (directional couplers, line stretchers, attenuators, etc.) feature extremely low flicker compared semiconductors and active devices, and (ii) the amplifier PM flicker is reduced proportionally to the carrier rejection because flicker comes from up-conversion of the near-dc noise.

The bridge exhibits low white noise as well. Neglecting trivial losses, the background noise is $b_0 = 2FkT_0/P_0$, where F is the amplifier noise factor, kT_0 is the thermal energy at room temperature, and P_0 is the power at the DUT output. The factor ‘2’ means that only half the DUT noise goes to the amplifier, the other half goes to the Σ port, but it can be reduced using an asymmetrical power combiner.

Finally, low sensitivity to AC magnetic fields from the power grid is easily achieved because RF amplification (20-40 dB) takes place before down-converting to DC.

No commercial solutions exist, but the scheme of Fig. 16.8 can be built on top of a commercial phase noise analyzer, including the modern dual-channel instruments.

The main ideas derive from Sann [Sann, 1968], but the RF amplification was introduced later by Labaar et al. [Labaar, 1982]. The method knew a sudden popularity in the late 1990s in Australia [Ivanov et al., 1998a, Ivanov et al., 1998b] and France [Rubiola et al., 1999]. We used it for the measurement of frequency stability in piezoelectric quartz resonators [Rubiola et al., 2000], and of the phase noise of DACs and DDSs [Calosso et al., 2020]. We demonstrated ultimate sensitivity adding the correlation for lowest white noise [Rubiola and Giordano, 2000], and a multi-stage bridge for lowest flicker [Rubiola and Giordano, 2002].

16.8.8 The Digital Phase-Noise and Allan-Variance Analyzer

Fig. 16.9 shows the block diagram of a digital phase-noise analyzer. The architecture differs from Fig. 9.1 in the use of Software Defined Radio (SDR) techniques. The input signal is digitized and down-converted to an I-Q stream at zero or near-zero frequency by multiplying the input data with sinusoidal signals from a NCO (Numerical Controlled Oscillator). The CORDIC algorithm [Volder, 1959, Volder, 2000, Meher et al., 2009] is probably the best option to convert the IQ stream into polar coordinates, phase and amplitude. The reference input signals A and B cannot be used as the sampling signal because the ADCs do not work well at arbitrary clock frequency. Consequently, the measurement of $\varphi - \theta$ requires two ADCs clocked by a common-mode oscillator, whose fluctuation is rejected. In turn, four ADCs are needed to reject the instrument noise.

The digital architecture has following interesting features

- It is great at measuring large values of $\varphi(t)$, not limited to $\pm\pi$, and even millions of cycles are not a problem.
- The use of a separate NCO at each input removes the requirement that the inputs are at the same frequency because all phases can be referred to the same ν_0 after a trivial numerical conversion.

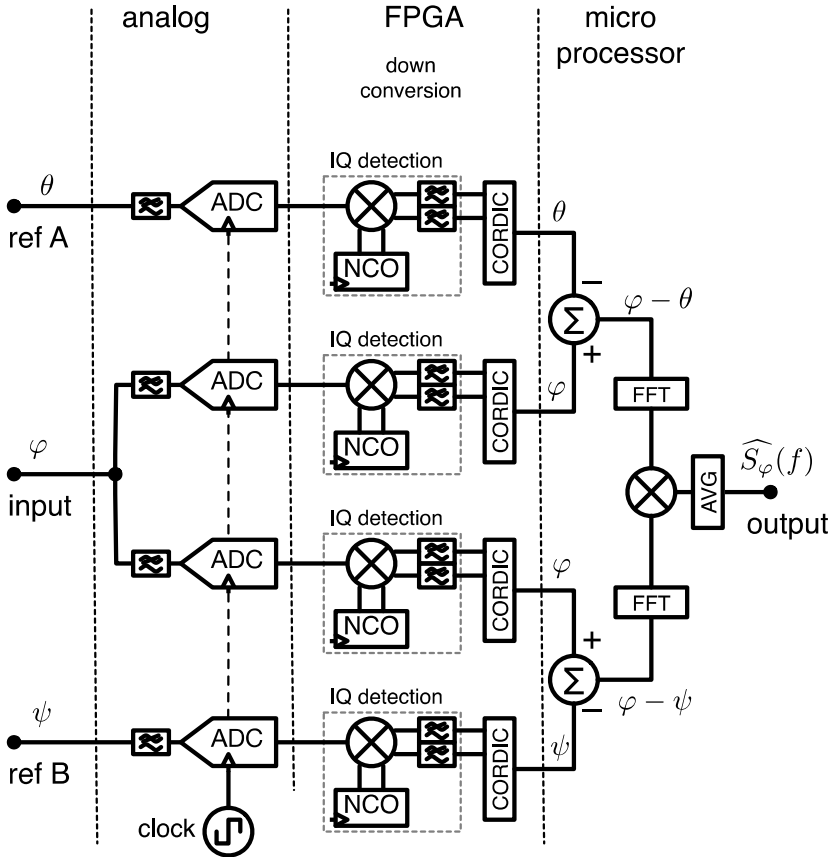


Figure 16.9: Digital phase noise and Allan variance analyzer.

- Full control on f_H in the measurement of Allan variances.

These properties make the digital architecture great at measuring both phase noise and variances, and open new perspectives.

Reference [Grove et al., 2004] pioneered the measurement of phase noise and the Allan variance with direct digitization of the RF signal. Later, [Mochizuki et al., 2007] provides a more detailed treatise, and [Sherman and Jördens, 2016] focuses on SDR techniques. The experimental methods for the noise characterization of ADCs for AM/PM noise measurements are shown in [Cárdenas-Olaya et al., 2017], and [Calosso et al., 2020] introduces a new method for the measurement of AM and PM noise in DACs and DDSs based on the amplification of the (random) modulation index, with optional AM/PM and PM/AM conversion.

Digital instruments are commercially available from Arcale (NoiseXT) and Microchip (Jackson Labs Technologies and Microsemi). The maximum frequency is limited to 30–400 MHz, depending on the instrument. Such instruments may have the reference oscillator(s) inside, or rely entirely on external oscillators. If the inputs

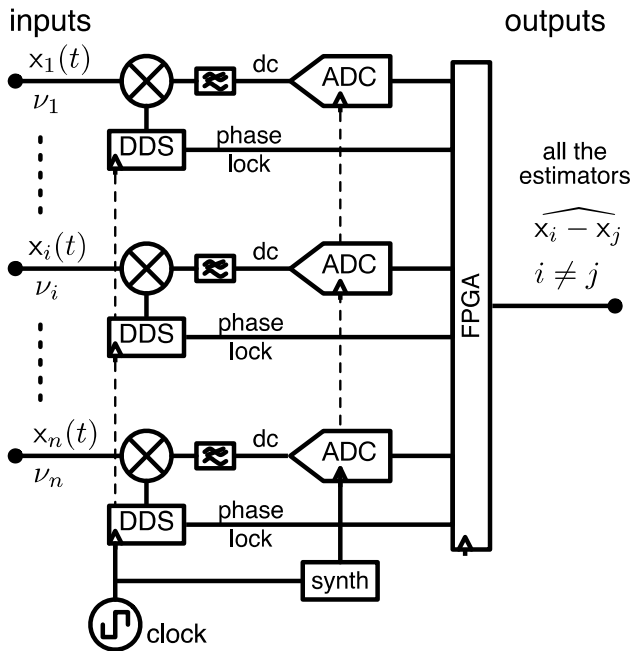


Figure 16.10: Tracking DDS.

A and B are accessed through a single connector instead of being available separately, the instrument cannot reject the noise of such oscillator. A hobby project is available from Andrew Holme's home page⁷, based on commercial boards, which even includes the source code and the binaries.

The Rohde & Schwarz analyzers FSWP and FSPN are based on similar concepts, but they implement microwave-to-IF down-conversion to extend the input range to 8, 26 or 50 GHz, depending on the model. They include two OCXOs and synthesizers as the references, and work with external references as well. The internal architecture of the FSWP is explained in [Feldhaus and Roth, 2016].

16.8.9 Tracking DDS (TDDS)

The TDDS [Calosso, 2013] is a phase locked loop where the voltage-controlled oscillator is replaced with a Direct Digital Synthesizer (DDS). The lock differs from a regular PLL in that the FPGA acts on the numerical phase of the DDS, instead of on the frequency of a voltage-controlled oscillator (VCO). Thus, an additional integrator is needed in the control loop. Interestingly, the TDDS locks in a wider frequency range than most PLLs, being limited by the frequency range of the mixer and of the DDS. The background noise is limited by the DDS (see [Calosso et al., 2012] for the DDS noise), noisier than double-balanced mixers.

⁷<http://www.aholme.co.uk/PhaseNoise/Main.htm>

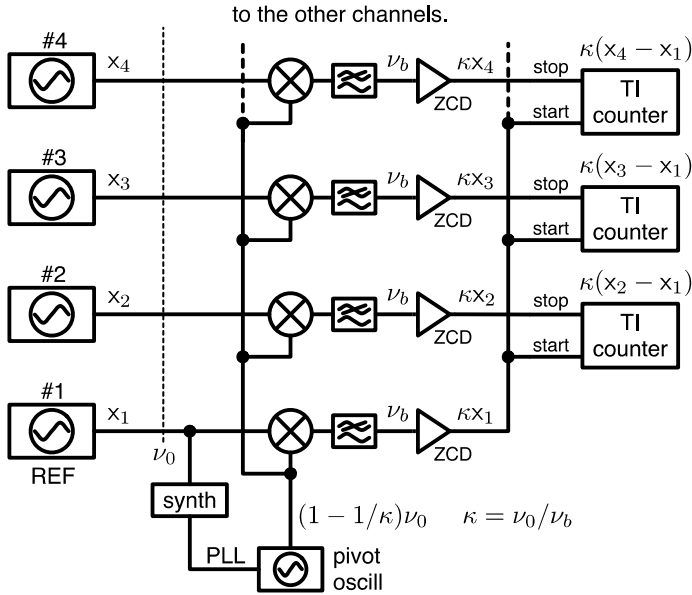


Figure 16.11: Multichannel analog Allan variance analyzer.

Fig. 16.10 shows a multichannel TDDS. At the start of operation, the FPGA sets the numerical frequency of all the DDSs to the frequency of the respective input. Then, the FPGA phase-locks each DDS to the input. At the same time, the FPGA provides the estimation of all the phase-time differences $\widehat{x_i - x_j}$ by combining the numerical phases and dc errors.

A multichannel TDDS was successfully used for the direct measurement of the frequency stability at the 100 MHz output of cryogenic sapphire oscillators [Calosso et al., 2019]. The Italian institute of metrology INRiM has developed the *Time Processor* a multichannel TDDS for the continuous monitoring and stability measurement of high-end oscillators and atomic frequency standards. Albeit some prototypes have already been transferred under contract to qualified users (European Space Agency and the FEMTO-ST Institute), until now there is only a conference publication [Calosso, 2021]. The Principal Investigator informed us⁸ that a startup Company is being created, under the provisional name Kairos TeX. The PicoPak⁹, a single channel TDDS, was produced by Precision Time and Frequency, LLC, but it is unsure whether it is still available.

16.8.10 Multichannel Analog Allan Variance Analyzer

This type of analyzer is a highly specialized instrument intended to simultaneously monitor multiple oscillators in a time scale (Cs beams, fountains, and H masers),

⁸Claudio E. Calosso, private communication, October 2022.

⁹Hamilton Technical Services, <http://www.wiley.com/7510A.pdf>.

usually comparing the 10 MHz or the 100 MHz outputs. Fig. 16.11 shows an example, easily extended to more than 4 inputs.

The machine exploits the leverage effect seen in Section 16.8.5, but in this case the pivot frequency ν_p is just below ν_0 , so that ν_b is in the sub-audio or low audio range. Accordingly, $x(t) = \varphi(t)/2\pi\nu$ is stretched by a factor $\kappa = \nu_0/\nu_b$. With $\kappa = 10^6$, a 100 MHz signal can be measured with 10 fs resolution using a simple counter that has a resolution of 10 ns. However, the actual resolution is limited by other factors.

The dual-mixer method, later extended to the multichannel system shown in Fig. 16.11, is introduced in [Allan and Daams, 1975]. Reference [Brida, 2002] discusses the design of the dual-mixer system, providing interesting experimental data. The best choice for the zero-crossing detector (ZCD) is a multistage amplifier where the first stage has narrow bandwidth for low noise, and the bandwidth increases progressively towards the output to allow high slew rate in the saturated signal [Dick et al., 1990, Collins, 1996].

Major metrology labs have been using instruments based on this principle for many years. However, those instruments are house-built prototypes. Commercial units are available from Quartzlock (UK), Time Tech (Germany), and VREMYA-CH (Russia).

Software Tools

Bibliography

- [ETS, 2002] (1998-2002). *Transmission and Multiplexing (TM); Generic requirements for synchronization networks*. European Telecommunications Standards Institute (ETSI), Sofia Antipolis, France. ETSI 300 462 Series, available from <https://etsi.org>.
- [ITU, 2017] (2017). *Recommendation ITU-R TF.538-4, Measures for random instabilities in frequency and time (phase)*. International Telecommunication Union (ITU), Geneva, CH.
- [Abdel Hafiz et al., 2022] Abdel Hafiz, M., Carlé, C., Passilly, N., Danet, J.-M., and Boudot, R. (2022). Light-shift mitigation in a microcell atomic clock with Symmetric Auto-Balanced Ramsey spectroscopy. *Appl. Phys. Lett.*, 120(6):064101.
- [Agilent-AN1449, 2003] Agilent-AN1449 (2003). *Fundamentals of RF and Microwave Power Measurements, Part 1–4*. Agilent Technologies, Inc., Paloalto, CA.
- [Allan, 1966] Allan, D. W. (1966). Statistics of atomic frequency standards. *Proc. IEEE*, 54(2):221–230.
- [Allan and Barnes, 1981] Allan, D. W. and Barnes, J. A. (1981). A modified “Allan variance” with increased oscillator characterization ability. In *Proc. Ann. Freq. Control Symp.*, pages 470–474, Ft. Monmouth, NJ.
- [Allan and Daams, 1975] Allan, D. W. and Daams, H. (1975). Picosecond time difference measurement system. In *Proc. Ann. Freq. Control Symp.*, pages 404–411, Atlantic City, NJ.
- [Allan and Levine, 2016] Allan, D. W. and Levine, J. (2016). A historical perspective on the development of the Allan variances and their strengths and weaknesses. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4).
- [Amelot et al., 2012] Amelot, J., Anand, D., Nelson, T., Stenbakken, G., Li-Baboud1, Y.-S., and Moyne, J. (2012). Towards timely intelligence in the power grid. In *Proc. Precision Time and Time Interval Conf.*, Reston, VA, USA.
- [Analog-MAT03,] Analog-MAT03. *Specification of the MAT-03 Low Noise Matched Dual PNP Transistor*. Analog Devices (formerly Precision Monolithics Inc.). Also available as mat03.pdf on <http://www.analog.com/>.
- [Babitch and Oliverio, 1974] Babitch, D. and Oliverio, J. (1974). Phase noise of oscillators at very low Fourier frequencies. In *Proc. Ann. Freq. Control Symp.*, pages 150–159, Atlantic City, NJ, USA.

- [Ball et al., 2016] Ball, H., Oliver, W. D., and Biercuk, M. J. (2016). The role of master clock stability in quantum information processing. *Nature P. J. Quantum Inf.*, 2:16033 1–8.
- [Banerjee, 2017] Banerjee, D. (2017). *PLL Performance, Simulation, and Design*. Dog Ear, 5 edition.
- [Barnes et al., 2011] Barnes, C. A., Hati, A., Nelson, C. W., and Howe, D. A. (2011). Residual PM noise evaluation of radio frequency mixers. In *Proc. Europ. Freq. Time Forum and Int’l Freq. Control Symp. Joint Meeting*, San Francisco, CA, USA.
- [Barnes, 1969] Barnes, J. A. (1969). *Tables of bias functions, B1 and B2, for variances based on finite samples of processes with power law spectral densities*, volume 375. US Department of Commerce, National Bureau of Standards.
- [Barnes et al., 1971] Barnes, J. A., Chi, A. R., Cutler, L. S., Healey, D. J., Leeson, D. B., McGunigal, T. E., Mullen, Jr, J. A., Smith, W. L., Sydnor, R. L., Vessot, R. F. C., and Winkler, G. M. R. (1971). Characterization of frequency stability. *IEEE Trans. Instrum. Meas.*, 20(2):105–120.
- [Barton and Leonov, 1998] Barton, D. K. and Leonov, S. A., editors (1998). *Radar Technology Encyclopedia*. Artech House, Nordwood, MA, USA.
- [Baudiquez et al., 2020] Baudiquez, A., Lantz, E., Rubiola, E., and Vernotte, F. (2020). Cross-spectrum measurement statistics: Uncertainties and detection limit. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 67(11):2461–2470.
- [Baudiquez et al., 2022] Baudiquez, A., Lantz, E., Rubiola, E., and Vernotte, F. (2022). The statistics of the cross-spectrum and the spectrum average: Generalization to multiple instruments. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 68(8):2585–2594.
- [Benkler et al., 2015] Benkler, E., Lisdat, C., and Sterr, U. (2015). On the relation between uncertainties of weighted frequency averages and the various types of Allan deviations. *Metrologia*, 55(4):565–574.
- [Bennett, 1948] Bennett, W. R. (1948). Spectra of quantized signals. *Bell Syst. Techn. J.*, 27(4):446–472.
- [Bernier, 1987] Bernier, L. G. (1987). Theoretical analysis of the modified Allan variance. In *Proc. Ann. Freq. Control Symp.*, pages 116–121, Philadelphia, PA, USA.
- [Blair, 1974] Blair, B. E., editor (1974). *Time and Frequency: Theory and Fundamentals*. US Department of Commerce. NBS Monograph 140.
- [Boudot and Rubiola, 2012] Boudot, R. and Rubiola, E. (2012). Phase noise in RF and microwave amplifiers. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 59(12):2613–2624.

- [Bregni, 2002] Bregni, S. (2002). *Synchronization of Digital Telecommunications Networks*. Wiley, Chichester, UK.
- [Bregni, 2016] Bregni, S. (2016). Twenty-five years of applications of the modified allan variance in telecommunications. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4).
- [Brida, 2002] Brida, G. (2002). High resolution frequency stability measurement system. *Rev. Sci. Instrum.*, 73(5):2171–2174.
- [Brigham, 1988] Brigham, O. E. (1988). *The Fast Fourier Transform and its Applications*. Prentice-Hall.
- [Bureau International des Poids et Mesures (BIPM), 2019] Bureau International des Poids et Mesures (BIPM) (2019). Le système international d’unités / the international system of units (‘the SI brochure’). Sèvres, France. Free available from the BIPM web site <https://bipm.org>.
- [Calosso, 2013] Calosso, C. E. (2013). Tracking DDS in time and frequency metrology. In *Proc. Europ. Freq. Time Forum and Int’l Freq. Control Symp. Joint Meeting*, pages 747–749, Prague, Czech Republic.
- [Calosso, 2021] Calosso, C. E. (2021). The time processor: A new platform for next generation timescales. In *ATIS’ Workshop on Synchronization and Timing Systems*, Vancouver, BC, Canada.
- [Calosso et al., 2016] Calosso, C. E., Clivati, C., and Micalizio, S. (2016). Avoiding aliasing in Allan variance: An application to fiber link data analysis. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4):646–655.
- [Calosso et al., 2020] Calosso, C. E., Cárdenas Olaya, A. C., and Rubiola, E. (2020). Phase-noise and amplitude-noise measurement of DACs and DDSs. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 67(2).
- [Calosso et al., 2012] Calosso, C. E., Gruson, Y., and Rubiola, E. (2012). Phase noise and amplitude noise in DDS. In *Proc. Int’l Freq. Control Symp.*, pages 777–782, Baltimore, MD, USA.
- [Calosso and Rubiola, 2017] Calosso, C. E. and Rubiola, E. (2017). Phase noise and jitter in digital electronics. arXiv:1701.00094 [physics.ins-det].
- [Calosso et al., 2019] Calosso, C. E., Vernotte, F., Giordano, V., Fluhr, C., Dubois, B., and Rubiola, E. (2019). Frequency stability measurement of cryogenic sapphire oscillators with a multichannel tracking DDS and the two-sample covariance. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 66(3).
- [Chi, 1966] Chi, A. R., editor (1966). *Proceedings IEEE vol. 54 no. 2, Special Issue on Frequency Stability*. IEEE.

- [Chi, 1964] Chi, A. R., editor (Nov. 23–24, 1964). *Short Term Frequency Stability, Proc. IEEE-NASA Symp. on the Definition and Measurement of Short-Term Frequency Stability*, volume NASA SP-80, Goddard Space Flight Center, Greenbelt, MD.
- [Christensen and Pearson, 1936] Christensen, C. J. and Pearson, G. L. (1936). Spontaneous resistance fluctuations in carbon microphones and other granular resistances. *Bell Syst. Techn. J.*, 15(2):197–223.
- [Collins, 1996] Collins, O. (1996). The design of low jitter hard limiters. *IEEE Trans. Commun.*, 44(5):601–608.
- [Conte et al., 2021] Conte, F., Massucco, S., Paolone, S., Schiapparelli, G.-P., Silvestro, F., and Zuo, Y. (2021). Frequency stability assessment of modern power systems: models definition and parameters identification. arXiv:2104.07330 [eess.SY].
- [Cundiff and Ye, 2003] Cundiff, S. T. and Ye, J. (2003). Colloquium: Femtosecond optical frequency combs. *Rev. Mod. Phys.*, 75(1):325–342.
- [Cárdenas-Olaya et al., 2017] Cárdenas-Olaya, A. C., Rubiola, E., Friedt, J.-M., Bourgeois, P.-Y., Ortolano, M., Micalizio, S., and Calosso, C. E. (2017). Noise characterization of analog to digital converters for amplitude and phase noise measurements. *Rev. Sci. Instrum.*, 88:065108 1–9.
- [Da Dalt and Sheikoleslami, 2018] Da Dalt, N. and Sheikoleslami, A. (2018). *Understanding Jitter and Phase Noise*. Cambridge University Press.
- [Dawkins et al., 2007] Dawkins, S. T., McFerran, J. J., and Luiten, A. N. (2007). Considerations on the measurement of the stability of oscillators with frequency counters. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 54(5):918–925.
- [Delva et al., 2017] Delva, P., Lodewyck, J., Bilicki, S., Bookjans, E., Vallet, G., Le Targat, R., Pottie, P.-E., Guerlin, C., Meynadier, F., Le Poncin-Lafitte, C., Lopez, O., Amy-Klein, A., Lee, W.-K., Quintin, N., Lisdar, C., Al-Masoudi, A., Dörscher, S., Grebing, C., Grosche, G., Kuhl, A., Raupach, S., Sterr, U., Hill, I. R., Hobson, R., Bowden, W., Kronjäger, J., Marra, G., Rolland, A., Baynes, F. N., Margolis, H. S., and Gill, P. (2017). Test of special relativity using a fiber network of optical clocks. *Phys. Rev. Lett.*, 118:221102 1–6.
- [Demir et al., 2000] Demir, A., Mehrotra, A., and Roychowdhury, J. (2000). Phase noise in oscillators: A unifying theory and numerical methods for characterization. *IEEE Trans. Circuits Syst. I. Fundam. Theory Appl. (1993-2003)*, 47(5):655–674.
- [Dennard et al., 1974] Dennard, R. H., Gaensslen, F. H., Rideout, V. L., Bassous, E., and LeBlanc, A. R. (1974). Design of ion-implanted mosfet’s with very small physical dimensions. *IEEE J. Solid-State Circuits*, 9(5):256–268. Also Proc. IEEE 87(4), Apr 1999.
- [Dick et al., 1990] Dick, G. J., Kuhnle, P. F., and Snyder, R. L. (1990). Zero-crossing detector with sub-microsecond jitter and crosstalk. In *Proc. Precision Time and Time Interval Conf.*, pages 269–282, Vienna, VI, USA.

- [Diddams et al., 2020] Diddams, S. A., Vahala, K., and Udem, T. (2020). Optical frequency combs: Coherently uniting the electromagnetic spectrum. *Science*, 369(6501):eaay3676.
- [Donley et al., 2022] Donley, (chair), E., Tavella, (vice-chair), P., Ashby, N., Calosso, C. E., Danielson, M., Galleani, L., Howe, D. A., Kinali, A., Koppang, P., Levine, J., Lin, C., Lutwak, R., Matsakis, D., Mitchell, P., Partridge, A., Radnaev, A., Rovera, D., Rubiola, E., Taylor, J., Vernotte, F., Vig, J., and Weaver, G. (2022). *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*. IEEE, New York.
- [Driscoll and Hanson, 1993] Driscoll, M. M. and Hanson, W. P. (1993). Measured vs. volume model-predicted flicker-of-frequency instability in VHF quartz crystal resonators. In *Proc. Int'l Freq. Control Symp.*, pages 186–192.
- [Edson, 1960] Edson, W. A. (1960). Noise in oscillators. *Proc. IRE*, 48(8):1454–1466.
- [Egan, 1990] Egan, W. F. (1990). Modeling phase noise in frequency dividers. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 37(4):307–315.
- [Egan, 2008] Egan, W. F. (2008). *Phase-Lock Basics*. Wiley & IEEE Press, New York, 2 edition.
- [Enzer et al., 2021] Enzer, D. G., Murphy, D. W., and Burt, E. A. (2021). Allan deviation of atomic clock frequency corrections: A new diagnostic tool for characterizing clock disturbances. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 68(7):2590–2601.
- [Esman et al., 2016] Esman, D., Ataie, V., Kuo, B. P.-P., Alic, N., and Radic, S. (2016). Subnoise signal detection and communication. *J. Lightw. Technol.*, 34(22):5214–5219.
- [Feldhaus and Roth, 2016] Feldhaus, G. and Roth, A. (2016). A 1 MHz to 50 GHz direct down-conversion phase noise analyzer with cross-correlation. In *Proc. Europ. Freq. Time Forum*, York, UK.
- [Fisher, 1935] Fisher, R. A. (1935). The fiducial argument in statistical inference. *Annals of Eugenics*, 5(4):391–398.
- [Frigo and Johnson, 2005] Frigo, M. and Johnson, S. G. (2005). The design and implementation of FFTW3. *Proc. IEEE*, 93(2):216–231.
- [Friis, 1944] Friis, H. T. (1944). Noise figure of radio receivers. *Proc. IRE*, 32:419–422.
- [Gabelli et al., 2004] Gabelli, J., Reydellet, L.-H., Fève, G., Berroir, J.-M., Plaçais, B., Roche, P., and Glatthli, D. C. (2004). Hanbury Brown – Twiss correlations to probe the population statistics of GHz photons emitted by conductors. *Phys. Rev. Lett.*, 93(5):056801.
- [Galleani, 2010] Galleani, L. (2010). The dynamic Allan variance II: A fast computational algorithm. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 57(1):182–188.

- [Galleani, 2011] Galleani, L. (2011). The dynamic Allan variance III: Confidence and detection surfaces. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 58(8):1550–1558.
- [Galleani and Tavella, 2009] Galleani, L. and Tavella, P. (2009). The dynamic Allan variance. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 56(3):450–464.
- [Galleani and Tavella, 2015] Galleani, L. and Tavella, P. (2015). The dynamic Allan variance IV: Characterization of atomic clock anomalies. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 62(5):791–801.
- [Galleani and Tavella, 2016] Galleani, L. and Tavella, P. (2016). The dynamic Allan variance V: Recent advances in dynamic stability analysis. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4):624–635.
- [Goldberg, 1999] Goldberg, B.-G. (1999). *Digital Frequency Synthesis Demystified*. LLH Technology Publishing, Eagle Rock, VA, USA.
- [Grebekemper, 1981] Grebekemper, C. J. (1981). Local oscillator phase noise and its effect on receiver performance. Technical Report Tech-notes vol. 8 no. 6, Watkins Johnson, San Jose, CA, USA.
- [Greenhall, 1998] Greenhall, C. A. (1998). Spectral ambiguity of Allan variance. *IEEE Trans. Instrum. Meas.*, 47(3):623–627.
- [Greenhall et al., 1999] Greenhall, C. A., Howe, D. A., and Percival, D. B. (1999). Total variance, an estimator of long-term frequency stability. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 46(5):1183–1191.
- [Greenhall and Riley, 2003] Greenhall, C. A. and Riley, W. J. (2003). Uncertainty of stability variances based on finite differences. In *Proc. Precision Time and Time Interval Conf.*, pages 267–280, San Diego, CA, USA.
- [Grove et al., 2004] Grove, J., Hein, J., Retta, J., Schweiger, P., Solbrig, W., and Stein, S. R. (2004). Direct-digital phase-noise measurement. In *Proc. Int’l Freq. Control Symp.*, pages 287–291, Montreal, Quebec, Canada.
- [Gruson et al., 2017] Gruson, Y., Giordano, V., Rohde, U. L., Poddar, A. K., and Rubiola, E. (2017). Cross-spectrum PM noise measurement, thermal energy, and metamaterial filters. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 64(3):634–642.
- [Gruson et al., 2020] Gruson, Y., Rus, A., Rohde, U. L., Roth, A., and Rubiola, E. (2020). Artifacts and errors in cross-spectrum phase noise measurements. *Metrologia*, 57(5):Art. no. 055010 p. 1–12. Open access.
- [Hajimiri and Lee, 1998] Hajimiri, A. and Lee, T. H. (1998). A general theory of phase noise in electrical oscillators. *IEEE J. Solid-State Circuits*, 33(2):179–194. Errata corrigé in vol. 33 no. 6 p. 928, June 1999.

- [Hajimiri and Lee, 1999] Hajimiri, A. and Lee, T. H. (1999). Design issues in CMOS differential LC oscillators. *IEEE J. Solid-State Circuits*, 34(5):717–724.
- [Halford et al., 1968] Halford, D., Wainwright, A. E., and Barnes, J. A. (1968). Flicker noise of phase in RF amplifiers: Characterization, cause, and cure. In *Proc. Ann. Freq. Control Symp.*, pages 340–341. (Abstract only is published).
- [Hall, 1960] Hall, R. N. (1960). Tunnel diodes. *IRE Trans. Electron. Dev.*, (?) (9):1–9.
- [Hanbury Brown et al., 1952] Hanbury Brown, R., Jennison, R. C., and Das Gupta, M. K. (1952). Apparent angular sizes of discrete radio sources. *Nature*, 170(4338):1061–1063.
- [Hanbury Brown and Twiss, 1956] Hanbury Brown, R. and Twiss, R. Q. (1956). Correlation between photons in two coherent beams of light. *Nature*, 177:27–29.
- [Harris, 1978] Harris, F. J. (1978). On the use of windows for harmonic analysis with the discrete fourier transform. *Proc. IEEE*, 66(1):172–204.
- [Hasegawa et al., 2011] Hasegawa, M., Chutani, R. K., Gorecki, C., Boudot, R., Dz-iuban, P., Giordano, V., Clatot, S., and Mauri, L. (2011). Microfabrication of Cesium vapor cells with buffer gas for MEMS atomic clocks. *Sensors Actuators A*, 67:594–601.
- [Hati et al., 2013] Hati, A., Nelson, C. W., Barnes, C., Lirette, D., Fortier, T., Quinlan, F., DeSalvo, J. A., Ludlow, A., Diddams, S. A., and Howe, D. A. (2013). State-of-the-art RF signal generation from optical frequency division. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 60(9):1796–1803.
- [Hellwig, 1988] Hellwig, (chair), H. (1988). *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology (IEEE Standard 1139-1988)*. IEEE, New York.
- [Holzer, 2018] Holzer, B., editor (2018). *Beam Injection, Extraction and Transfer*. CERN. Licence CC-BY 4.0.
- [Horn, 1969] Horn, C. H. (1969). A carrier suppression technique for measuring S/N and carrier/sideband ratios greater than 120 dB. In *Proc. Ann. Freq. Control Symp.*, pages 223–235, Fort Monmouth, NJ.
- [Howe, 2000] Howe, D. A. (2000). The total deviation approach to long-term characterization of frequency stability. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 47(5):1102–1110.
- [Howe, 2006] Howe, D. A. (2006). Thêoh: a hybrid, high-confidence statistic that improves on the allan deviation. *Metrologia*, 43(4):S322–S331.
- [Howe et al., 2000] Howe, D. A., Beard, R., Greenhall, C., Vernotte, F., and Riley, B. (2000). The total estimator of the Hadamard function used for GPS operations. In *Proc. Precision Time and Time Interval Conf.*

- [Howe and Peppler, 2003] Howe, D. A. and Peppler, T. K. (2003). Very long-term frequency stability: Estimation using a special-purpose statistic. In *Proc. Intl. Freq. Control Symp. and Europ. Freq. Time Forum Joint Meeting*, pages 233–238, Tampa, FL, USA.
- [Huang, 2015] Huang, A. B. (2015). Moore’s law is dying (and that could be good). *IEEE Spectr.*, pages 41–44.
- [ITU Working Group 15, 2000] ITU Working Group 15 (2000). *Definitions and terminology for synchronization in packet networks, Recommendation G.8260*. International Telecommunication Union, Geneva, CH.
- [Ivanov et al., 1998a] Ivanov, E. N., Tobar, M. E., and Woode, R. A. (1998a). Application of interferometric signal processing to phase-noise reduction in microwave oscillators. *IEEE Trans. Microw. Theory Techn.*, 46(10):1537–1545.
- [Ivanov et al., 1998b] Ivanov, E. N., Tobar, M. E., and Woode, R. A. (1998b). Microwave interferometry: Application to precision measurements and noise reduction techniques. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 45(11):1526–1535.
- [Jablonski et al., 2015] Jablonski, S., Schlarb, H., and Sydlo, C. (2015). CW laser based phase reference distribution for particle accelerators. In *Proc. Int’l Beam Instrumentation Conf. (IBIC2015)*, Melbourne, Australia.
- [Johansson, 2005] Johansson, S. (2005). New frequency counting principle improves resolution. In *Proc. Ann. Freq. Control Symp.*, pages 628–635, Vancouver, BC, Canada.
- [Johnson, 1928] Johnson, J. B. (1928). Thermal agitation of electricity in conductors. *Phys. Rev.*, 32:97–109.
- [Joint Committee for Guides in Metrology (JCGM), 2012] Joint Committee for Guides in Metrology (JCGM) (2012). *International Vocabulary of Metrology — Basic and General Concepts and Associated Terms (VIM) document JCGM 200:2012*, 3 edition. (The VIM is available free of charge. The 4th edition of the VIM, named VIM4, is in progress.).
- [Jones et al., 2003] Jones, D. J., Holman, K. W., Notcutt, M., Ye, J., Chandalia, J., Jiang, L. A., Ippen, E. P., and Yokoyama, H. (2003). Ultralow-jitter, 1550-nm mode-locked semiconductor laser synchronized to a visible optical frequency standard. *Optics Lett.*, 28(10):813–815.
- [Kalistz, 2004] Kalistz, J. (2004). Review of methods for time interval measurements with picosecond resolution. *Metrologia*, 41:17–32.
- [Kester, 2004] Kester, W., editor (2004). *Analog-Digital Conversion*. Analog Devices, USA.
- [Kihara, 1989] Kihara, M. (1989). Performance aspects of reference clock distribution for evolving digital networks. *IEEE Communications Mag.*, 27(4):24–34.

- [Kramer and Klische, 2001] Kramer, G. and Klische, W. (2001). Multi-channel synchronous digital phase recorder. In *Proc. Int'l Freq. Control Symp.*, pages 144–151, Seattle, WA, USA.
- [Kramer and Klische, 2004] Kramer, G. and Klische, W. (2004). Extra high precision digital phase recorder. In *Proc. Europ. Freq. Time Forum*, pages 595–602, Guildford, UK.
- [Krieger and Younis, 2006] Krieger, G. and Younis, M. (2006). Impact of oscillator noise in bistatic and multistatic SAR. *Geosci. Remote Sens. Lett.*, 3(3):424–428.
- [Kroupa, 1983] Kroupa, V. F., editor (1983). *Frequency Stability: Fundamentals and Measurement*. IEEE Press, New York.
- [Kroupa, 1998] Kroupa, V. F. (1998). The state of the art of flicker frequency noise in BAW and SAW quartz resonators. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 35(3):406–420.
- [Kroupa, 1999] Kroupa, V. F., editor (1999). *Direct Digital Frequency Synthesizers*. IEEE Press, New York.
- [Kroupa, 2001] Kroupa, V. F. (2001). Jitter and phase noise in frequency dividers. *IEEE Trans. Instrum. Meas.*, 50(5):1241–1243.
- [Kroupa, 2005] Kroupa, V. F. (2005). Theory of $1/f$ noise—a new approach. *Phys. Lett. A*, (336):126–132.
- [Kroupa, 2012] Kroupa, V. F. (2012). *Frequency Stability*. IEEE Wiley.
- [Kruse et al., 2021] Kruse, J., Schäfer, B., and Witthaut, D. (2021). Revealing drivers and risks for power grid frequency stability with explainable ai. *Patterns*, 2(11):100365.
- [Kuse and Fermann, 2017] Kuse, N. and Fermann, M. E. (2017). Electro-optic comb based real time ultra-high sensitivity phase noise measurement system for high frequency microwaves. *Scientific Reports*, 7(1):2847, 1–8.
- [Kärtner, 1990] Kärtner, F. X. (1990). Analysis of white and $f^{-\alpha}$ noise in oscillators. *Int. J. Circ. Theor. Appl.*, 18:485–519.
- [Labaar, 1982] Labaar, F. (1982). New discriminator boosts phase noise testing. *Microwaves*, 21(3):65–69.
- [Lee and Hajimiri, 2000] Lee, T. H. and Hajimiri, A. (2000). Oscillator phase noise: A tutorial. *IEEE J. Solid-State Circuits*, 35(3):326–336.
- [Leeson, 1966] Leeson, D. B. (1966). A simple model of feed back oscillator noise spectrum. *Proc. IEEE*, 54:329–330.
- [Leeson, 2016] Leeson, D. B. (2016). Oscillator phase noise: A 50-year review. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(8):1208–1225. Open Access.

- [Leeson and Johnson, 1966] Leeson, D. B. and Johnson, G. F. (1966). Short-term stability for a Doppler radar: Requirements, measurements and techniques. *Proc. IEEE*, 54(2):244–248.
- [Leeson and Johnson, 1964] Leeson, D. B. and Johnson, G. F. (Nov. 23–24, 1964). Short-term stability for a doppler radar: Requirements, measurements and techniques. In [Chi, 1964], pages 3–9.
- [Lesage and Audoin, 1973] Lesage, P. and Audoin, C. (1973). Characterization of frequency stability: uncertainty due to the finite number of measurements. *IEEE Trans. Instrum. Meas.*, 22(2):157–161. See also corrections published in IM-23 p. 103, IM-24 p. 86, and IM-25 p. 270.
- [Levine et al., 2016] Levine, J., Tavella, P., and Santarelli, G., editors (2016). *IEEE Transact. Ultrason. Ferroelec. Frequency Control, Special Issue on Celebrating the 50th Anniversary of the Allan Variance*. IEEE.
- [Levitin et al., 2012] Levitin, D. J., Chordiac, P., and Menon, V. (2012). Musical rhythm spectra from Bach to Joplin obey a $1/f$ power law. *Proc. USA Nat. Acad. of Sci.*, 109(10).
- [L(f)-2014, 2014] L(f)-2014 (2014). European Workshop on Cross-Spectrum Phase Noise Measurements. Laboratoire National de métrologie et d’Essays (LNE), Paris, France.
- [L(f)-2015, 2015] L(f)-2015 (2015). Cross Spectrum $\mathcal{L}(f)$ Workshop. Denver, CO, USA. Side event at the 2015 Europ. Freq. Time Forum and Freq. Control Symp. Joint Meeting, April 13–16.
- [L(f)-2017, 2017] L(f)-2017 (2017). Cross Spectrum $\mathcal{L}(f)$ Workshop. Besançon, France. Side event at the 2017 Europ. Freq. Time Forum and Freq. Control Symp. Joint Meeting, July 10–13.
- [Li, 2008] Li, M. P. (2008). *Jitter, Noise, and Signal Integrity at High-Speed*. Prentice Hall, Boston, MA, USA.
- [Lin, 2018] Lin, J.-C. (2018). Synchronization requirements for 5G. *IEEE Veh. Technol. Mag.*, 13(3):91–99.
- [Lindley, 1958] Lindley, D. V. (1958). Fiducial distributions and bayes’ theorem. *J. Royal Statistical Soc. Series B (Methodological)*, 20(1):102–107.
- [Lisdat et al., 2016] Lisdat, C., Grosche, G., Quintin, N., Shi, C., Raupach, S. M. F., Grebing, C., Nicolodi, D., Stefani, F., Al-Masoudi, A., Dörscher, S., Häfner, S., Robyr, J.-L., Chiodo, N., Bilicki, S., Bookjans, E., Koczwara, A., Koke, S., Kuhl, A., Wiotte, F., Meynadier, F., Camisard, E., Abgrall, M., Lours, M., Legero, T., Schnatz, H., Sterr, U., Denker, H., Chardonnet, C., Le Coq, Y., Santarelli, G., Amy-Klein, A., Le Targat, R., Lodewyck, J., Lopez, O., and Pottier, P.-E. (2016). A clock network for geodesy and fundamental science. *Nature Communications*, 7(1):12443.

- [Loh et al., 2013a] Loh, W., Yegnanarayanan, S., Ram, R. J., and Juodawlkis, P. W. (2013a). Unified theory of oscillator phase noise I: White noise. *IEEE Trans. Microw. Theory Techn.*, 61(6):2371–2381.
- [Loh et al., 2013b] Loh, W., Yegnanarayanan, S., Ram, R. J., and Juodawlkis, P. W. (2013b). Unified theory of oscillator phase noise II: Flicker noise. *IEEE Trans. Microw. Theory Techn.*, 61(12):4130–4144.
- [Mack, 2015] Mack, C. (2015). The multiple lives of Moore’s law. *IEEE Spectr.*, pages 29–35.
- [Malkin, 2016] Malkin, Z. (2016). Application of the allan variance to time series analysis in astrometry and geodesy: A review. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4).
- [Meher et al., 2009] Meher, P. K., Valls, J., Tso-Bing, J., Sridharan, K., and Maharatna, K. (2009). 50 years of CORDIC: Algorithms, architectures, and applications. *IEEE Trans. Circuits Syst. I, Reg. Papers*, 56(9):1893–1907.
- [Milotti, 2002] Milotti, E. (2002). $1/f$ noise: a pedagogical review.
- [Mochizuki et al., 2007] Mochizuki, K., Uchino, M., and Morikawa, T. (2007). Frequency-stability measurement system using high-speed ADCs and digital signal processing. *IEEE Trans. Instrum. Meas.*, 56(5):1887–1893.
- [Moore’s Law, 2015] Moore’s Law (2015). Moore’s law 50 years. A series of articles and editorials on IEEE Spectrum.
- [Narbonneau et al., 2003] Narbonneau, F., Lours, M., Lopez, O., Clairon, A., and Santarelli, G. (2003). Ultra-stable optical links for metrological applications. In *Proc. Freq. Control Symp. and European Freq. and Time Forum Joint Meeting*, pages 1041–1047, Tampa, FL.
- [Nelson et al., 2014] Nelson, C. W., Hati, A., and Howe, D. A. (2014). A collapse of the cross-spectral function in phase noise metrology. *Rev. Sci. Instrum.*, 85(3):024705, 1–7.
- [Nicolodi et al., 2014] Nicolodi, D., Argence, B., Zhang, W., Le Targat, R., Santarelli, G., and Le Coq, Y. (2014). Spectral purity transfer between optical wavelengths at the 10^{-18} level. *Nature Photonics*, 8(3):219–223.
- [Numata et al., 2004] Numata, K., Kemery, A., and Camp, J. (2004). Thermal-noise limit in the frequency stabilization of lasers with rigid cavities. *Phys. Rev. Lett.*, pages 250602 1–4.
- [Nyquist, 1928] Nyquist, H. (1928). Thermal agitation of electric charge in conductors. *Phys. Rev.*, 32:110–113.
- [Oliver, 1965] Oliver, B. M. (1965). Thermal and quantum noise. *Proc. IEEE*, 53(5):436–454.

- [Owen, 2004] Owen, D. (2004). Good practice to phase noise measurement. Teddington, UK.
- [Pankratz and Sánchez-Sinencio, 2014] Pankratz, E. and Sánchez-Sinencio, E. (2014). Survey of integrated-circuit-oscillator phase-noise analysis. *Int. J. Circ. Theor. Appl.*, 42(9):871–938.
- [Penfield, 1960] Penfield, Jr., P. (1960). Noise in negative-resistance amplifiers. *IRE Trans. Circuit Theory*, 7(2):166–170.
- [Phillips, 1987] Phillips, D. (1987). Random noise in digital gates and dividers. In *Proc. Int'l Freq. Control Symp.*, pages 507–511, Philadelphia, PA, USA.
- [Razavi, 2012] Razavi, B. (2012). *RF Microelectronics*. Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition.
- [Reinhardt, 2005] Reinhardt, V. S. (2005). A review of time jitter and digital systems. In *Proc. Int'l Freq. Control Symp.*, pages 38–45.
- [RG-CE System Protection and Dynamics Sub Group, 2016] RG-CE System Protection and Dynamics Sub Group (2016). Frequency stability evaluation criteria for the synchronous zone of continental europe — requirements and impacting factors. Technical report, European Network of Transmission System Operators for Electricity (ENTSOE), Brussels, Belgium.
- [Riehle, 2015] Riehle, F., editor (2015). *Proc. 8th Frequency Standards and Metrology Symp.*, Potsdam, Germany. IOP. Published as vol. 723, 2016 of Journal of Physics: Conference Series.
- [Riley,] Riley, W. J. Handbook of frequency stability analysis. NIST Special Publication 1065.
- [Riley and Greenhall., 2004] Riley, W. J. and Greenhall., C. A. (2004). Power law noise identification using the lag 1 autocorrelation. In *Proc. Europ. Freq. Time Forum*, Guildford, UK.
- [Robins, 1984] Robins, W. P. (1984). *Phase Noise in Signal Sources*. Peter Peregrinus and IEE.
- [Rohde et al., 2021] Rohde, U. L., Rubiola, E., and Whitaker, J. C. (2021). *Microwave and wireless synthesizers*. John Wiley & Sons.
- [Rubiola, 2005a] Rubiola, E. (2005a). The measurement of AM noise of oscillators. arXiv:physics/0512082.
- [Rubiola, 2005b] Rubiola, E. (2005b). On the measurement of frequency and of its sample variance with high-resolution counters. *Rev. Sci. Instrum.*, 76(5). Also arXiv:physics/0411227, Dec. 2004.
- [Rubiola, 2006] Rubiola, E. (2006). Tutorial on the double-balanced mixer. <http://arxiv.org, document arXiv:physics/0608211>.

- [Rubiola, 2010] Rubiola, E. (2010). *Phase Noise and Frequency Stability in Oscillators*. Cambridge Univ. Press, Cambridge, UK.
- [Rubiola and Boudot, 2007] Rubiola, E. and Boudot, R. (2007). The effect of AM noise on correlation phase noise measurements. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 54(5):926–932. Also arXiv:physics/0609147.
- [Rubiola and Brendel, 2010] Rubiola, E. and Brendel, R. (2010). A generalization of the Leeson effect. arXiv:1003.0113 [physics.ins-det].
- [Rubiola and Giordano, 2000] Rubiola, E. and Giordano, V. (2000). Very high frequency and microwave interferometric PM and AM noise measurements. *Rev. Sci. Instrum.*, 71(8):3085–3091.
- [Rubiola and Giordano, 2002] Rubiola, E. and Giordano, V. (2002). Advanced interferometric phase and amplitude noise measurements. *Rev. Sci. Instrum.*, 73(6):2445–2457.
- [Rubiola and Giordano, 2007] Rubiola, E. and Giordano, V. (2007). On the $1/f$ frequency noise in ultra-stable quartz oscillators. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 54(1):15–22. Also arXiv:physics/0602110.
- [Rubiola et al., 1999] Rubiola, E., Giordano, V., and Gros Lambert, J. (1999). Very high frequency and microwave interferometric PM and AM noise measurements. *Rev. Sci. Instrum.*, 70(1):220–225.
- [Rubiola et al., 2000] Rubiola, E., Giordano, V., and Gros Lambert, J. (2000). Flicker noise measurement of HF quartz resonators. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 47(2):361–368.
- [Rubiola et al., 2016] Rubiola, E., Lenczner, M., Bourgeois, P.-Y., and Vernotte, F. (2016). The Ω counter, a frequency counter based on the linear regression. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(7):961–969. Special Issue of the 2015 IFCS EFTF joint meeting.
- [Rubiola et al., 2005] Rubiola, E., Salik, E., Huang, S., and Maleki, L. (2005). Photonic delay technique for phase noise measurement of microwave oscillators. *J. Opt. Soc. Am. B - Opt. Phys.*, 22(5):987–997.
- [Rubiola and Vernotte, 2010] Rubiola, E. and Vernotte, F. (2010). The cross-spectrum experimental method. arXiv:1004.5539 [physics.ins-det].
- [Rutman, 1978] Rutman, J. (1978). Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress. *Proc. IEEE*, 66(9):1048–1075.
- [Rutman and Walls, 1991] Rutman, J. and Walls, F. L. (1991). Characterization of frequency stability in precision frequency sources. *Proc. IEEE*, 79(6):952–960.

- [Salik et al., 2004] Salik, E., Yu, N., Maleki, L., and Rubiola, E. (2004). Dual photonic-delay-line cross correlation method for the measurement of microwave oscillator phase noise. In *Proc. Europ. Freq. Time Forum and Freq. Control Symp. Joint Meeting*, pages 303–306, Montreal, Canada.
- [Sann, 1968] Sann, K. H. (1968). The measurement of near-carrier noise in microwave amplifiers. *IEEE Trans. Microw. Theory Techn.*, 9:761–766.
- [Saulson, 1990] Saulson, P. R. (1990). Thermal noise in mechanical experiments. *Phys. Rev. D*, 42(8).
- [Schottky, 1926] Schottky, W. (1926). Small-shot effect and flicker effect. *Phys. Rev.*, 28(74):74–103.
- [Serrano et al., 2011] Serrano, J., Alvarez, P., Lipinski, M., and Wlostowski, T. (2011). Accelerator timing system overview. In *Proc. Particle Accelerator Conf. (PAC’11)*, New York, NY, USA.
- [Sherman and Jördens, 2016] Sherman, J. A. and Jördens, R. (2016). Oscillator metrology with software defined radio. *Rev. Sci. Instrum.*, 87(5):054711.
- [Shoaf, 1974] Shoaf, J. H. (1974). Specification and measurement of frequency stability. NBS Technical Notes 74-396, NBS (now NIST).
- [Shu and Sánchez-Sinencio, 2005] Shu, K. and Sánchez-Sinencio, E. (2005). *CMOS PLL Synthesizers*. Springer, Boston, MA.
- [Sikula, 2003] Sikula, J., editor (2003). *Noise and Fluctuations*. CNRL, Brno, Czech Republic. Proc. 2003 Intl. Conf. on Noise and 1/f Fluctuations.
- [Skolnik, 2001] Skolnik, M. I. (2001). *Introduction to Radar Systems*. McGraw Hill, New York, NY, USA, 3 edition.
- [Skolnik, 2008] Skolnik, M. I., editor (2008). *Radar Handbook*. McGraw Hill, New York, NY, USA, 3 edition.
- [Sosin, 1971] Sosin, B. M. (1971). H.F. communication receiver performance requirements and realization. *The Radio and Electronic Engineer*, 41(7):321–329.
- [Stein, 2010] Stein, S. R. (2010). The Allan variance—challenges and opportunities. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 57(3):540–547.
- [Sthal et al., 2013] Sthal, F., Devel, M., Ghosh, S., Imbaud, J., Cibieli, G., and Bourquin, R. (2013). Volume dependence in handel’s model of quartz crystal resonator noise. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 60(9):1971–1977.
- [Sullivan et al., 1990] Sullivan, D. B., Allan, D. W., Howe, D. A., and L., W. F., editors (1990). *Characterization of Clock and Oscillators, NIST Technical Note 1337*. NIST.

- [Tataria et al., 2021] Tataria, H., Shafi, M., Molisch, A. E., Dohler, M., Sjöland, H., and Tufvesson, F. (2021). 6G wireless systems: Vision requirements, challenges, insights, and opportunities. *Proc. IEEE*, 109(7):1166–1199.
- [Taylor and Howe, 2010] Taylor, J. A. and Howe, D. A. (2010). Fast theobor: A method for long data set analysis. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 57(9):2091–2094.
- [Udem et al., 2002] Udem, T., Holzwarth, R., and Hänsch, T. W. (2002). Optical frequency metrology. *Nature*, 416(6877):233–237.
- [van der Pol and van der Mark, 1927] van der Pol, B. and van der Mark, J. (1927). Frequency demultiplication. *Nature*, 120:363–364.
- [van der Ziel, 1962] van der Ziel, A. (1962). On the noise figure of negative-conductance amplifiers. *IRE Trans. Circuit Theory*, 9(1):83–84.
- [Verbruggen et al., 1989] Verbruggen, A. H., Stoll, H., Heeck, K., and Koch, R. H. (1989). A novel technique for measuring resistance fluctuations independently of background noise. *Appl. Phys. A*, 48:233–236.
- [Vernotte et al., 2021] Vernotte, F., Chen, S., and Rubiola, E. (2021). Response and uncertainty of the parabolic variance PVAR to noninteger exponents of power law. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 70.
- [Vernotte and Lantz, 2012] Vernotte, F. and Lantz, E. (2012). Statistical biases and very long term time stability analysis. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 59(3):523–530.
- [Vernotte et al., 2016] Vernotte, F., Lenczner, M., Bourgeois, P.-Y., and Rubiola, E. (2016). The parabolic variance (PVAR), a wavelet variance based on the least-square fit. *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, 63(4):611–623. Special Issue of the 50th anniversary of the Allan variance.
- [Vernotte et al., 1998a] Vernotte, F., Zalamansky, G., and Lantz, E. (1998a). Time stability characterization and spectral aliasing. Part I: a time-domain approach. *Metrologia*, 35(5):723–730.
- [Vernotte et al., 1998b] Vernotte, F., Zalamansky, G., and Lantz, E. (1998b). Time stability characterization and spectral aliasing. Part II: a frequency-domain approach. *Metrologia*, 35(5):731–738.
- [Vessot et al., 1964] Vessot, R. F. C., Mueller, R. F., and Vanier, J. (1964). A cross-correlation technique for measuring the short-term properties of stable oscillators. In *Short Term Frequency Stability, Proc. IEEE-NASA Symp. on the Definition and Measurement of Short-Term Frequency Stability*, pages 231–235.
- [Victor, 1956] Victor, W. K. (1956). The evaluation of phase-stable oscillators for coherent communication systems. In *Proc. Ann. Freq. Control Symp.*, pages 268–304, Asbury Park, NJ, USA.

- [Volder, 1959] Volder, J. E. (1959). The CORDIC trigonometric computing technique. *IRE Trans. Electron. Comput.*, 8(3):330–334.
- [Volder, 2000] Volder, J. E. (2000). The birth of CORDIC. *JVLSISP*, 25(2):101–105.
- [Volyanskiy et al., 2008] Volyanskiy, K., Cussey, J., Tavernier, H., Salzenstein, P., Sauvage, G., Larger, L., and Rubiola, E. (2008). Applications of the optical fiber to the generation and to the measurement of low-phase-noise microwave signals. *J. Opt. Soc. Am. B - Opt. Phys.*, 25(12):2140–2150. Also arXiv:0807.3494v1 [physics.optics].
- [Wallin, 2018] Wallin, A. E. E. (2018). Power law noise identification for allantools. <https://ui.adsabs.harvard.edu/abs/2018ascl.soft04021W>.
- [Walls et al., 1992] Walls, F. L., Handel, P. H., Besson, R., and Gagnepain, J.-J. (1992). A new model of $1/f$ noise in baw quartz resonators. In *Proc. Int'l Freq. Control Symp.*, pages 327–333.
- [Walls et al., 1976] Walls, F. L., Stein, S. R., Gray, J. E., and Glaze, D. J. (1976). Design considerations in state-of-the-art signal processing and phase noise measurement systems. In *Proc. Ann. Freq. Control Symp.*, pages 269–274, Atlantic City, NJ, USA.
- [Walter, 1994] Walter, T. (1994). Characterizing frequency stability: A continuous power-law model with discrete sampling. *IEEE Trans. Instrum. Meas.*, 43(1):69–79.
- [Welch, 1967] Welch, P. D. (1967). The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.*, 15(2):70–73.
- [Weste and Harris, 2011] Weste, N. H. E. and Harris, D. M. (2011). *CMOS VLSI Design, A Circuits and Systems Perspective*. Addison Wesley, Boston, MA, USA, 4 edition.
- [Widrow and Kollár, 2008] Widrow, B. and Kollár, I. (2008). *Quantization Noise*. Cambridge Univ. Press.
- [Widrow et al., 1996] Widrow, B., Kollár, I., and Liu, M.-C. (1996). Statistical theory of quantization. *IEEE Trans. Instrum. Meas.*, 45(2):353–361.
- [Xie et al., 2017] Xie, X., Bouchand, R., Nicolodi, D., Giunta, M., Hänsel, W., Lezius, M., Joshi, A., Datta, S., Alexandre, C., Lours, M., Tremblin, P.-A., Santarelli, G., Holzwarth, R., and Le Coq, Y. (2017). Photonic microwave signals with zeptosecond-level absolute timing noise. *Nature Photonics*, 11(1):44–47.

Notation, Symbols and Acronyms

Notation and symbols

dot (\dot{a}) derivative over time

bar (\bar{a}) mean, or weighted average

hat (\hat{a}) estimation

\leftrightarrow Fourier or Laplace transform inverse-transform pair

$\langle \rangle$ mathematical expectation, same as $\mathbb{E}\{ \}$

$\langle \rangle_\tau$ average over a specified time interval τ

$\langle \rangle_n$ average over an integer number n of samples

B bandwidth of the baseband signal, and of AM/PM noise (the bandwidth of the RF signal is $2B$)

b_n coefficients of the polynomial law of $S_\varphi(f)$

D_y fractional-frequency drift

\mathfrak{D} degrees of freedom (often denoted with ν , dof or DOF in the literature about statistics)

d_n coefficients of the polynomial law of $S_\nu(f)$

$\mathbb{E}\{ \}$ mathematical expectation

F noise factor (see NF below)

$\mathcal{F}\{ \}$ Fourier transform

f Fourier frequency, in spectral analysis

f_c corner frequency, where flicker noise equals white noise

f_L resonator's Leeson frequency (baseband bandwidth), half of the RF bandwidth

h_n coefficients of the polynomial law of $S_y(f)$

$H(f)$ transfer function, typically in $|H(f)|^2$

k_n coefficients of the polynomial law of $S_x(f)$

$\mathcal{L}\{ \}$ Laplace transform

$\mathcal{L}(f)$ phase noise, usually $10 \log_{10}[\mathcal{L}(f)]$ [dBc/Hz]

N white noise PSD [W/Hz]

NF Noise Figure, $NF = 10 \log_{10}(F)$

P carrier power

Q quality factor, in resonators

$S_x(f)$ one-sided power spectral density of the random variable $x(t)$. For us, chiefly α , φ , x , y and ν

- T time interval, or period
 T acquisition time (data record for one FFT)
 T (subscript) signal truncated over a duration T , as in $x_T(t)$ and $X_T(f)$
 \mathcal{T} total acquisition time (full data record used in the measurement)
 t time
 V_0 peak amplitude (of the clock signal)
 $w(t; \tau)$ weight function (counters), or wavelet-like function (variances)
 $x(t)$ time fluctuation
 $y(t)$ fractional frequency fluctuation
 $\alpha(t)$ fractional amplitude fluctuation
 $(\Delta\nu)(t)$ frequency fluctuation
 $\epsilon(t)$ amplitude fluctuation
 $\theta(t)$ random phase, replacement for $\varphi(t)$, when needed
 Λ triangular average, or a frequency counter implementing triangular average
 ν carrier frequency
 Π uniform average, or a frequency counter implementing uniform average
 $\sigma_x^2(\tau)$ same as TVAR, used in formulas
 $^A\sigma_y^2(\tau)$ same as AVAR, used in formulas
 $^M\sigma_y^2(\tau)$ same as MVAR, used in formulas
 $^P\sigma_y^2(\tau)$ same as PVAR, used in formulas
 $^H\sigma_y^2(\tau)$ same as HVAR, used in formulas
 τ delay (delay line)
 τ measurement (integration) time
 τ relaxation time, in resonators
 τ time shift (correlation, convolution, etc.)
 $\varphi(t)$ random phase
 $\psi(t)$ random phase, replacement for $\varphi(t)$, when needed
 ω shorthand for $2\pi\nu$
 Ω shorthand for $2\pi f$
 Ω parabolic-weight average, or a frequency counter implementing such average

Acronyms

ADC Analog to Digital Converter

AM Amplitude Modulation (also AM noise)

ADEV Square root of AVAR

AVAR Allan VARiance

BER Bit Error Rate (digital electronics and telecom)

BPF Band Pass Filter

CORDIC COordinate Rotation Digital Computer, or Volder's algorithm

DAC Digital to Analog Converter

DBM Double Balanced Mixer

DDS Direct Digital Synthesizer

DLL Delay Locked Loop

DUT Device Under Test (oscillator or two-port component)

FFT Fast Fourier Transform

FLL Frequency Locked Loop

FM Frequency Modulation (also FM noise)

FPGA Field Programmable Gate Array

FS FemtoSecond (laser)

GRADEV Gap Resistant ADEV

HDEV Square root of HVAR

HVAR Hadamard VARiance

IF Intermediate Frequency (superheterodyne receiver). Also IF output of a mixer

IQ In-phase and Quadrature (detection, modulation)

LO Local Oscillator (superheterodyne). Also LO input of a mixer

LPF Low Pass Filter

LSB Least significant bit, in signal processing and digital electronics

LSB Lower Side Band, in modulated signals

MDEV Square root of MVAR

MSB Most significant bit, in signal processing and digital electronics

NCO Numerically Controlled Oscillator

MVAR Modified [Allan] VARiance

PLL Phase Locked Loop

PM Phase Modulation (also PM noise)

PDF Probability Density Function

PSD Power Spectral Density

PDEV Square root of PVAR

PVAR Parabolic VARiance

RF Radio Frequency. Also RF input of a mixer

SDR Software Defined Radio

TDDS Tracking Direct Digital Synthesizer

TIE Time Interval Error [ITU Working Group 15, 2000]

TVAR Time VARiance





USB Upper Side Band, in modulated signals

VCO Voltage Controlled Oscillator. Also voltage-control input of an oscillator

ZCD Zero Crossing Detector

Frequency Counter		Wavelet Variance	
$\bar{y}(\tau) = \int_{\mathbb{R}} y(t) w(t; \tau) dt$		$\sigma_y^2(\tau) = \mathbb{E} \left\{ \frac{1}{2} [\bar{y}_2 - \bar{y}_1]^2 \right\} = \mathbb{E} \left\{ \int_{\mathbb{R}} [y(t) w(t; \tau)]^2 dt \right\}$	
type of frequency counter	Π	associated variance	AVAR $A \sigma_y^2(\tau)$
	Λ		MVAR $M \sigma_y^2(\tau)$
	Ω		PVAR $P \sigma_y^2(\tau)$
	(Δ)		HVAR $H \sigma_y^2(\tau)$
Note: this representation is only for theoretical purposes. There are no commercial Δ counters.		Note: the wavelet representation hides the second difference in the evaluation of the variance.	

Other Two-Sample Variances



Spectra to Variances Conversion

noise type	$S_y(f)$	ANAR $A_{\sigma_y^2}(\tau)$	MVAR $M_{\sigma_y^2}(\tau)$	HVAR $H_{\sigma_y^2}(\tau)$	PVAR $P_{\sigma_y^2}(\tau)$	TVAR $T_{\sigma_y^2}(\tau)$
Blue PM	$h_3 f^3$	$\frac{3f_h^2 h_3}{8\pi^2 \tau^2}$ $0.0380 f_H h_3 / \tau^2$	$\frac{10\gamma + \ln 48 + 10 \ln(\pi f_H \tau)}{16\pi^4}$ $\frac{h_3}{16\pi^4}$ $[10\gamma + \ln 48 + 10 \ln \pi] = 0.0135$	$\frac{5f_h^2 h_3}{18\pi^2 \tau^2}$ $0.0281 f_H h_3 / \tau^2$	$\frac{9[\gamma + \ln(4\pi f_H \tau)]}{4\pi^4}$ $\frac{h_3}{4\pi^4} = 0.0718$	$\frac{10\gamma + \ln 48 + 10 \ln(\pi f_H \tau)}{48\pi^2}$ $\frac{h_3}{48\pi^2} = 0.00451$
White PM	$h_2 f^2$	$\frac{3f_H h_2}{4\pi^2 \tau^2}$ $0.0760 f_H h_2 / \tau^2$	$\frac{3 h_2}{8\pi^2 \tau^2}$ $0.0380 h_2 / \tau^2$	$\frac{5f_H h_2}{9\pi^2 \tau^2}$ $0.0563 f_H h_2 / \tau^2$	$\frac{3 h_2}{2\pi^2 \tau^2}$ $0.152 h_2 / \tau^2$	$\frac{1 h_2}{8\pi^2 \tau}$ $0.0127 h_2 / \tau$
Flicker PM	$h_1 f$	$\frac{3\gamma - \ln 2 + 3 \ln(2\pi f_H \tau)}{4\pi^2}$ $\frac{h_1}{\tau^2}$ $[3\gamma - \ln 2 + 3 \ln(2\pi f_H \tau)] / 4\pi^2 = 0.166$	$\frac{(24 \ln 2 - 9 \ln 3) h_1}{8\pi^2 \tau^2}$ $\frac{h_1}{\tau^2}$ $\simeq \frac{5[\gamma + \ln(\sqrt[3]{48} \pi f_H \tau)]}{9\pi^2}$ $\frac{h_1}{\tau^2}$ $5[\gamma + \ln(\sqrt[3]{48} \pi)] / 9\pi^2 = 0.119$	$\frac{3\gamma - \ln 2 + 3 \ln(2\pi f_H \tau)}{4\pi^2}$ $\frac{h_1}{\tau^2}$ $\simeq \frac{5[\gamma + \ln(\sqrt[3]{48} \pi f_H \tau)]}{9\pi^2}$ $\frac{h_1}{\tau^2}$ $5[\gamma + \ln(\sqrt[3]{48} \pi)] / 9\pi^2 = 0.119$	$\frac{3[\ln(16) - 1] h_1}{2\pi^2 \tau^2}$ $\frac{h_1}{\tau^2}$ $\frac{3 h_0}{5 \tau}$	$\frac{(8 \ln 2 - 3 \ln 3) h_1}{8\pi^2}$ $\frac{h_1}{0.0285 h_1}$
White FM	h_0	$\frac{1 h_0}{2 \tau}$ $0.935 h_{-1}$	$\frac{1 h_0}{4 \tau}$ $0.935 h_{-1}$	$\frac{1 h_0}{3 \tau}$ $0.750 h_{-1}$	$\frac{3 h_0}{5 \tau}$ $1.60 h_{-1}$	$\frac{1 h_0 \tau}{12}$ $0.312 h_{-1} \tau^2$
Flicker FM	$h_{-1} f^{-1}$	$\frac{2 \ln(2) h_{-1}}{1.39 h_{-1}}$ $0.935 h_{-1}$	$\frac{27 \ln 3 - 32 \ln 2}{8} h_{-1}$ $0.935 h_{-1}$	$\frac{8 \ln 2 - 3 \ln 3}{3} h_{-1}$ $0.750 h_{-1}$	$\frac{2[7 - \ln(16)] h_{-1}}{5}$ $1.60 h_{-1}$	$\frac{27 \ln 3 - 32 \ln 2}{24} h_{-1} \tau^2$ $0.312 h_{-1} \tau^2$
Random walk FM	$h_{-2} f^{-2}$	$\frac{2\pi^2}{3} h_{-2} \tau^2$ $5.43 h_{-2} \tau^2$	$\frac{11\pi^2}{20} h_{-2} \tau^2$ $5.43 h_{-2} \tau^2$	$\frac{2\pi^2}{9} h_{-2} \tau^2$ $2.19 h_{-2} \tau^2$	$\frac{26\pi^2}{35} h_{-2} \tau^2$ $7.33 h_{-2} \tau^2$	$\frac{11\pi^2}{60} h_{-2} \tau^3$ $1.81 h_{-2} \tau^3$
Integrated flicker FM	$h_{-3} f^{-3}$	not convergent	not convergent	$\frac{\pi^2[27 \ln 3 - 32 \ln 2]}{9} h_{-3} \tau^2$ $8.205 h_{-3} \tau^2$	not convergent	not convergent
Integrated RW FM	$h_{-4} f^{-4}$	not convergent	not convergent	$\frac{44\pi^2}{90} h_{-4} \tau^3$ $4.825 h_{-4} \tau^3$	not convergent	not convergent
linear drift D_y		$\frac{1}{2} D_y^2 \tau^2$	$\frac{1}{2} D_y^2 \tau^2$	0	$\frac{1}{2} D_y^2 \tau^2$	$\frac{1}{6} D_y^2 \tau^4$
Spectral response $ H(\theta) ^2, \theta = \pi f \tau$		$\frac{2 \sin^4 \theta}{\theta^2}$	$\frac{2 \sin^6 \theta}{\theta^4}$	$\frac{16 \sin^6 \theta}{9 \theta^2}$	$\frac{9[2 \sin^2 \theta - \theta \sin 2\theta]^2}{2 \theta^6}$	$\frac{\tau^2}{3} \frac{2 \sin^6 \theta}{\theta^4}$

$\gamma = 0.577$ is the Euler Mascheroni constant. Formulae hold for $\tau \gg f_H/2$ where appropriate, f_H = bandwidth (sharp cutoff filter). MVAR, PVAR and TVAR require $\tau \gg \tau_0$, where τ_0 = sampling interval.

About this book



Enrico Rubiola